Review: Disk Imaging Solutions p. 61

Windows IT Pro

A PENTON PUBLICATION

OCTOBER 2010 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

Kerberos Logons Revealed

p. 23

Advanced AD Security p. 28

Highly Available iSCSI Storage for VMware p. 32

Secure SSL Certificates p. 38

Exchange 2010 Messaging Records Management p. 41

Debugging PowerShell p. 45

Exchange Server Client Access: Load Balance Your Servers p. 48

Escape SharePoint Permissions Purgatory p. 54

Building the engines of a Smarter Planet:

Workloads will grow. Your number of servers doesn't have to.

As a smarter planet generates more and more data, the strains on existing infrastructures intensify. Yet in today's competitive business environment, midsize businesses need more performance with fewer resources. Fortunately, a solution now exists from IBM: the first x86 servers with memory decoupled from their processor, which eliminates the need to buy additional servers to support growing workloads. It's enterprise-level performance without enterprise budgets. Here's how IBM and our Business Partners can help you work smarter:

intel Xeon inside™

Powerful. Intelligent.

1 **Get up to 261% more performance.** The IBM System x3690 X5, featuring the Intel® Xeon® processor 7500 series, is the first scalable 2-socket system designed to offer the performance, memory capacity and reliability of an enterprise 4-socket system.[1]

2 **Simplify your IT.** The new x3690 can accomplish the same job as 33 servers[2] from previous generations. Consolidating your aging servers can simplify your IT while minimizing server sprawl, and can help reduce the costs of managing multiple systems.

3 **Gain 82% more virtual servers for the same license cost.**[3] Now you can virtualize to reduce complexity. With IBM's unique MAX5 memory that scales independently of the processor, you'll have the available memory you need for your demanding virtualization environment.

The IBM System x3690 X5 is designed and priced with midsize businesses in mind. Starting at $275 per month for 36 months.[4]

4 **Learn how you could see ROI in as little as 3 months.** With increased performance and simplified IT, you can dramatically reduce overall IT costs such as power, cooling and facilities by up to 98%.[2]

Midsize businesses are the engines of a Smarter Planet.

Connect with an IBM Business Partner now and start reaping the benefits of smarter systems. Call **1-877-IBM-ACCESS** or visit **ibm.com**/systems/x3690 to access the IBM Systems Consolidation Evaluation Tool.

IBM®

[1] Based on 3-party benchmark result SPEC CPU2006 comparing previous generation system with Intel Xeon 5470 (Harpertown) 3.33GHz processors to new generation System x3690 X5 with Intel Xeon X7560 processors. [2] Return on investment and power savings calculation based on 33:1 consolidation ratio scenario of 200 Intel 2-socket servers to 6 IBM x3690 X5 servers and savings in energy costs, software license fees and other operating costs. Actual costs and savings will vary depending on individual customer configurations and environment. [3] Based on one IBM x3690 X5 with MAX5 memory expansion (2 processors, 64 DIMMs) as compared to an industry-standard 2-processor, 18 DIMM system and license cost per processor. [4] Prices are current as of 8-31-10 and are subject to change without notice. Manufacturer's suggested retail price; dealer prices may vary. Minimum transaction size $5,000; monthly payments are estimates based on lease rates for installations of qualified products and services in the United States. Actual rates may vary based on your creditworthiness, configuration details, etc., and are subject to credit approval by IBM Credit LLC. For some clients, total software and services is limited to 75% of hardware financed. Other conditions may apply, so please contact your IBM Authorized Business Partner or IBM representative for more information. IBM, the IBM logo, ibm.com, Smarter Planet and the planet icon are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml. Intel, the Intel logo, Xeon and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

CONTENTS

COVER STORY

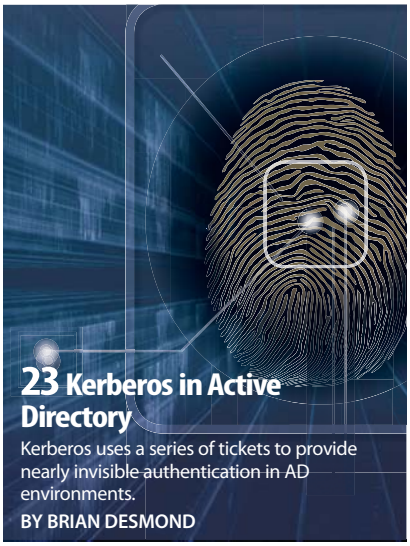

23 Kerberos in Active Directory

Kerberos uses a series of tickets to provide nearly invisible authentication in AD environments.

BY BRIAN DESMOND

FEATURES

28 Advanced Active Directory Security

Securing Active Directory can be tricky if you're not aware of protected groups—and though you might not ever have to change the default settings, it's useful to know how AdminSDHolder works.

BY RUSSELL SMITH

SOLUTIONS PLUS

32 Configuring Highly Available iSCSI Storage for VMware ESX Server 4.x

Walk through a sample scenario of creating a LUN that's mirrored between two SAN servers, configuring redundant network connections from the VMware ESX server to each of the SAN servers, and setting up an initiator and target for successful connectivity.

BY GREG SHIELDS

38 Are Your SSL Certificates Secure?

Follow these steps to determine key lengths of your current certificates and find out about infrastructure concerns when upgrading from certificates with 1,024-bit key lengths to 2,048-bit key lengths.

BY ALAN SUGANO

41 Exchange 2010 MRM: Implementing New Retention Policies

Instead of using managed folders like its predecessor, Exchange 2010's messaging records management (MRM) system uses retention tags and policies. Here's how to design, create, and apply retention tags and policies using PowerShell.

BY TONY REDMOND

45 Debugging in Windows PowerShell

Bugs are an inevitable part of life when you write PowerShell scripts. Here are some basic techniques for hunting down and squashing bugs and some practices that can help keep bugs at bay.

BY DON JONES

48 Exchange Server's Client Access: Load Balancing Your Servers

In Exchange Server 2010, Client Access servers handle all client connections, so it's important you configure them with load balancing. Learn about situations that call for establishing persistence and how to load balance HTTP and MAPI traffic.

BY KEN ST. CYR

54 Escaping SharePoint Permissions Purgatory

If you've recently upgraded from WSS 2.0 or SharePoint Portal Server 2003 to SharePoint 2007, you're probably battling permissions problems. Here's the solution.

BY RYAN THOMAS

INTERACT

17 Reader to Reader

Quickly identify the administrators and privileged users who recently logged on to a server from a remote machine.

19 Ask the Experts

Learn about virtualizing domain controllers, use jumbo frames in ESX, enable Hyper-V on a box without the required hardware, and create a shortcut to lock Windows.

IN EVERY ISSUE

6 IT Community Forum

79 Directory of Services

79 Advertising Index

79 Vendor Directory

80 Ctrl+Alt+Del

Access articles online at www.windowsitpro.com. Enter the article ID (located at the end of each article) in the InstantDoc ID text box on the home page.

Windows IT Pro
A PENTON PUBLICATION

OCTOBER 2010
VOLUME_16
NO_10

COLUMNS

CROCKETT | IT PRO PERSPECTIVES

5 Survey Confirms Past Year Hard on IT Budgets

Although IT organizations made drastic budget cuts in the past year, cloud computing and virtualization are likely to re-loosen the purse strings.

THURROTT | NEED TO KNOW

9 What You Need to Know About Small Business Server "Aurora," the Great Windows Phone 7 Debate, IE 9, and More

As Microsoft ships its upcoming Windows Small Business Server product, code-named Aurora, it also is upgrading Windows Home Server and Internet Explorer. Find out why we say these are all good releases.

MINASI | WINDOWS POWER TOOLS

12 ImageX Provides Disk Imaging on a Budget

This free command-line Ghost alternative lets you quickly copy disk images onto troubled PCs and restore them to service.

OTEY | TOP 10

14 Remote Desktop Keyboard Shortcuts

Your standard keyboard shortcuts won't work in a remote session; here are useful shortcuts to know for Remote Desktop sessions, such as how to switch programs, start Task Manager, and take screenshots of the remote session.

WHEELER | WHAT WOULD MICROSOFT SUPPORT DO?

15 An Easier Way to View Incoming WMI Queries

Troubleshooting performance problems related to Windows Management Instrumentation (WMI) can be difficult because of hard-to-read WMI tracing logs. But with help from a script created by the Microsoft support team, you can view WMI queries on your systems in a reader-friendly format.

NEW TECHNOLOGY!

All virtual machines suffer from fragmentation.

Until now, there was no solution.

V-locity™ 2

Virtual Platform Disk Optimizer
for VMware® and Hyper-V™

Ensure maximum I/O performance on every virtual O/S.

V-locity 2 gives you the management power you need.

- Optimize the performance on your entire virtualized platform from the host disk to the VMs
- Eliminate resource management priority conflicts
- Maximize I/O bandwidth efficiency
- Eliminate "bloated" free space on thin/dynamic disks
- Do it all automatically, in the background, with zero resource conflicts

Try it FREE for 30 days. No obligation.
www.diskeeper.com/v2

Or call now for a volume license discount quote: **800-829-6468**

Diskeeper corporation

C O N T E N T S

PRODUCTS

58 New & Improved
Check out the latest products to hit the marketplace.
PRODUCT SPOTLIGHT: Diskeeper Corporation's **V-locity 2.0**

REVIEW
59 Paul's Picks
Small Business Server "Aurora" is the right product at the right time—and forget cut and paste, Windows Phone has Xbox LIVE.
BY PAUL THURROTT

REVIEW
60 SpamTitan
SpamTitan is a software message filter that protects against spam and viruses with multiple technologies and granular policy controls in a simple UI for admins and end users.
BY NATHAN WINTERS

COMPARATIVE REVIEW
61 3 Disk Imaging Solutions —Redux
Since we reviewed these products in 2008, the market has changed but the basic concepts have remained the same. This time, we look at Acronis Snap Deploy, Paragon Deployment Manager, and Symantec Ghost with a focus on how they differ.
BY ERIC B. RUX

BUYER'S GUIDE
65 Power Management Software for Windows Workstations
One way to cut costs is to reduce workstations' energy consumption. With power management software, you can centrally configure and manage the power management settings for numerous workstations—including those running older client OSs—without writing any scripts.
BY KAREN BEMOWSKI

72 Industry Bytes
Learn how to multitask in PowerShell 2.0, see where we're at with Unified Communications and where it's going, discover what's next in VMware's vSphere, and find out how to avoid treating your networks like Han Solo.

Windows IT Pro

EDITORIAL

Editorial and Custom Strategy Director
Michele Crockett mcrockett@windowsitpro.com

Executive Editor, IT Group
Amy Eisenberg amy@windowsitpro.com

Technical Director
Michael Otey motey@windowsitpro.com

Senior Technical Analyst
Paul Thurrott news@windowsitpro.com

Custom Group Editorial Director
Dave Bernard dbernard@windowsitpro.com

Web and Developer Strategic Editor
Anne Grubb agrubb@windowsitpro.com

Systems Management
Karen Bemowski kbemowski@windowsitpro.com
Caroline Marwitz cmarwitz@windowsitpro.com
Zac Wiggy zwiggy@windowsitpro.com

Messaging , Mobility, SharePoint, and Office
Brian Keith Winstead bwinstead@windowsitpro.com

Networking and Hardware
Jason Bovberg jbovberg@windowsitpro.com

Security
Lavon Peters lpeters@windowsitpro.com

SQL Server
Megan Bearly Keller mkeller@windowsitpro.com
Sheila Molnar smolnar@windowsitpro.com

Editorial Web Architect
Brian Reinholz breinholz@windowsitpro.com

IT Media Group Editors
Linda Harty, Chris Maxcer, Rita-Lyn Sanders

CONTRIBUTORS

SharePoint and Office Community Editor
Dan Holme danh@intelliem.com

Senior Contributing Editors
David Chernicoff david@windowsitpro.com
Mark Joseph Edwards mje@windowsitpro.com
Kathy Ivens kivens@windowsitpro.com
Mark Minasi mark@minasi.com
Paul Robichaux paul@robichaux.net
Mark Russinovich mark@sysinternals.com

Contributing Editors
Alex K. Angelopoulos aka@mvps.org
Sean Deuby sdeuby@windowsitpro.com
Michael Dragone mike@mikerochip.com
Jeff Fellinge jeff@blackstatic.com
Brett Hill brett@iisanswers.com
Darren Mar-Elia dmarelia@windowsitpro.com
Tony Redmond 12knocksinna@gmail.com
Ed Roth eroth@windowsitpro.com
Eric B. Rux ericbrux@whshelp.com
John Savill john@savilltech.com
William Sheldon bsheldon@interknowlogy.com
Randy Franklin Smith rsmith@montereytechgroup.com
Curt Spanburgh cspanburgh@scg.net
Orin Thomas orin@windowsitpro.com
Douglas Toombs help@toombs.us
Ethan Wilansky ewilansky@windowsitpro.com

ART & PRODUCTION

Production Director
Linda Kirchgesler linda@windowsitpro.com

Senior Graphic Designer
Matt Wiebe matt.wiebe@penton.com

ADVERTISING SALES

Publisher
Peg Miller pmiller@windowsitpro.com

Director, International and Agency Services
Don Knox don.knox@penton.com

Business Development Director
Kerry Gates kerry.gates@penton.com

EMEA Managing Director
Irene Clapham irene.clapham@penton.com

Director of IT Strategy and Partner Alliances
Birdie J. Ghiglione birdie.ghiglione@penton.com
619-442-4064

Online Sales and Marketing Manager
Dina Baird Dina.Baird@penton.com

Key Account Director
Chrissy Ferraro christina.ferraro@penton.com
970-203-2883

Account Executives
Barbara Ritter barbara.ritter@penton.com
858-367-8058
Cass Schulz cassandra.schulz@penton.com
858-357-7649

Client Project Managers
Michelle Andrews 970-613-4964
Kim Eck 970-203-2953

Ad Production Supervisor
Glenda Vaught glenda.vaught@penton.com

MARKETING & CIRCULATION

Customer Service service@windowsitpro.com

IT Group Audience Development Director
Marie Evans marie.evans@penton.com

Marketing Director
Sandy Lang sandy.lang@penton.com

CORPORATE

PENTON

Chief Executive Officer
Sharon Rowlands Sharon.Rowlands@penton.com

Chief Financial Officer/Executive Vice President
Jean Clifton jean.clifton@penton.com

TECHNOLOGY GROUP

Senior Vice President, Technology Media Group
Kim Paulsen kpaulsen@windowsitpro.com

WRITING FOR *WINDOWS IT PRO*
Submit queries about topics of importance to Windows managers and systems administrators to articles@windowsitpro.com.

LIST RENTALS
Contact MeritDirect, 333 Westchester Avenue, White Plains, NY or www.meritdirect.com/penton.

REPRINTS
Diane Madzelonka, Diane.madzelonka@penton.com, 216-931-9268, 888-858-8851

Yet Another 10 Free Tools for System Administrators

Audit Active Directory and file servers, detect inactive users, block USB devices, and more – for free

The following freeware tools by Windows IT Pro Community Choice Awards finalist NetWrix Corporation can save you a lot of time and make your network more efficient – at absolutely no cost. Some of these tools have advanced commercial versions with additional features, but none of them will expire and stop working when you urgently need them.

10. Disk Space Monitor (MS TechNet Magazine Sep'09: www.tinyurl.com/23kt6lt) — Even with today's terabyte-large hard drives, server disk space tends to run out quickly and unexpectedly. This simple monitoring tool will send you daily reports regarding all servers that are running low on disk space, below the configurable threshold. Download link: www.tinyurl.com/29hp8c5

9. Bulk Password Reset (reviewed by SoftPedia: www.tinyurl.com/27bmaj7) – While most companies have strong password policies for their employees, one critical issue is still neglected: local Administrator passwords on all servers are usually managed in a "set and forget" fashion, sometimes using some "well-known" passwords, opening a major surface for security attacks. The Bulk Password Reset tool quickly resets local account passwords on all servers at once, making them more secure. Download link: www.tinyurl.com/2cmaolj

8. Windows Service Monitor (WindowsReference.com: www.tinyurl.com/28yx29o) —This very simple monitoring tool alerts you when some Windows service accidentally stops on one of your servers. The tool also detects services that fail to start at boot time, which sometimes happens, for example, with Exchange Server. Download link: www.tinyurl.com/2ex8hod

7. VMware Change Reporter (TechTarget/SearchVirtualDesktop: www.tinyurl.com/2cujnrt) — If you don't know what is being changed by your colleagues in the VMware infrastructure, it's very easy to get lost and miss changes that can affect the things for which you are responsible. This tool tracks and reports configuration changes in VMware Virtual Center settings and permissions. Download link: www.tinyurl.com/28j6vjp

6. Active Directory Object Restore Wizard (4sysops.com: www.tinyurl.com/23gflz3) — This tool can save the day if someone accidentally (or intentionally) deleted a bunch of Active Directory objects. It provides granular object-level and even attribute-level restore capabilities to quickly rollback unwanted changes (e.g., mistakenly deleted users, modified group memberships, etc). Download link: www.tinyurl.com/24hc3mz

5. File Server Change Reporter (4sysops.com: www.tinyurl.com/2eb2d6u) — This tool enhances the line of auditing tools; this one for file servers. File Server Change Reporter detects changes in files, folders, permissions, tracks deleted, and newly-created files, and sends daily summary reports. This is a very useful tool to detect mistakenly-deleted files and recover from backup or to see if someone changes some important files. Download link: www.tinyurl.com/2aoysb2

4. Inactive Users Tracker (MS TechNet Magazine May'08: www.tinyurl.com/2fycw5b) — This feature tracks down inactive user accounts (e.g., terminated employees) so you can easily disable them, or even remove them entirely, to eliminate potential security holes. The tool sends reports on a regular schedule, showing what accounts have been inactive for a configurable period of time (e.g., 2 months). Download link: www.tinyurl.com/2fglsl7

3. Password Expiration Notifier (Redmond Magazine Feb'09, 4sysops: www.tinyurl.com/258elbh) — This tool will automatically remind users to change passwords before they expire to keep you safe from password reset calls. It works nicely for users who don't log on interactively and, thus, never receive standard password change reminders at log on time (e.g., VPN and OWA users). Download: www.tinyurl.com/265kxb2

2. USB Blocker (Windows IT Pro Nov'09: InstantDoc ID 102860: www.tinyurl.com/24ls5zx) — Users bring tons of consumer devices: flash drives, MP3 players, cell phones, etc., into the office and this aptly-named tool can block them with a couple of mouse clicks to prevent the spread of a virus and to restrict the take-out of confidential information. The product is integrated with Active Directory and is very easy to use. Download link: www.tinyurl.com/2g95fn8

1. Active Directory Change Reporter (Windows IT Pro Sep'09: InstantDoc ID 102446, Windows IT Pro Jan'09: InstantDoc ID 100593, TechTarget: www.tinyurl.com/24sqv7c) — This is a simple auditing tool to keep tabs on what's going on inside Active Directory. The tool tracks changes to users, groups, OUs, and other types of AD objects, and sends summary reports with full lists of what was changed and how it was changed. In addition, it has a nice "rollback" feature that helps rollback unwanted changes (including deletions) very quickly. Download link: www.tinyurl.com/25fm6l7

Crockett

"The most alarming signal of the declining economy's effect on IT organizations was the cut in security and business continuity technology."

Survey Confirms Past Year Hard on IT Budgets

But amidst across-the-board IT cuts, new technology slips in

IT organizations took a serious whack at their budgets this past year, according to our annual survey of *Windows IT Pro* magazine readers conducted by independent publication research company Readex. The drop in gross revenues year over year among the companies responding sets the stage: Average gross revenue for IT organizations responding to our survey dropped from $4 billion in 2009 to $2.8 billion in 2010. (Keep in mind that among the respondents are IT pros who represent some very large companies.)

Staffing expenditures dropped from an average of $6.6 million in 2009 to $6.2 million in 2010, about a 6 percent drop—a minor decrease compared with other drops in spending. Spending on consulting and outsourcing was down from an average of $3.55 million in 2010 to $1.37 million in 2010—a staggering 61 percent decline. The drop in average spending on computer systems plummeted from $6.53 in 2009 to $2.33 million in 2010. Spending on software dropped by 57 percent. But the most alarming signal of the declining economy's effect on IT organizations was the cut in security and business continuity technology—down 58 percent from 2009, from $3.3 million in 2009 to $1.38 million in 2010.

> Organizations are gradually shifting to Windows Server 2008. About 49 percent are using Server 2008; 53 percent are using Server 2008 R2.

Some New Technologies Still Creeping In

Despite these grim numbers, some new technology is slipping into IT organizations and onto users' desktops (or laptops). About 66 percent of respondents reported that Windows 7 is in use in their organizations, and 48 percent are using Windows 7 x64 Edition. Windows Vista inevitably took the fall, down from 94 percent of respondents reporting its use in their organizations in 2009 to only 39 percent in 2010. Windows XP is exhibiting far more staying power—its use dropped only slightly: 94 percent of respondents reported that their organizations used XP in 2009, whereas 93 percent reported its use in 2010.

In the server category, our respondents reported that organizations are gradually shifting to Windows Server 2008. About 49 percent of organizations are using Windows Server 2008 and 53 percent are using Windows Server 2008 R2. Use of Windows Server 2003 is down from 91 percent in 2009 to 79 percent in 2010. Windows 2000 Server use dropped from 36 percent in 2009 to 21 percent in 2010. And Windows NT Server is still kicking, but use is down slightly from 6 percent in 2009 to 5 percent in 2010.

Server products competing with Microsoft technologies took a hit this year among our audience. Linux Server use is down from 46 percent in 2009 to 39 percent in 2010. HP-UX use is down from 13 percent in 2009 to 8 percent in 2010. And although it's no surprise that Exchange Server dominates among our audience in the messaging category, the adoption of Exchange Server 2010 continues to move at a glacial pace. In 2009, 9 percent of respondents were using Exchange Server 2010. In 2010, that number has crept up to 12 percent. Among virtualization products, Hyper-V use rose from 13 percent in 2009 to 17 percent in 2010.

One of the newest categories of spending that we're beginning to track with this year's survey is spending on cloud computing. By far the biggest areas of spending are in messaging, followed by collaboration services. Vertical-market spending is evenly split across a number of categories, including human resources, customer relationship management, financial services, sales management, and supply chain management.

Critical Technologies

Cloud computing and virtualization will be the most likely technologies to loosen organizations' budgets. After a year of serious slashing, organizations are beginning to see some declines in performance and services, which might hamper plans to seize opportunities as the economy finally recovers. Of course, no one yet knows when the real recovery will happen.

InstantDoc ID 125870

MICHELE CROCKETT (michele.crockett@penton.com) helped launch *SQL Server Magazine* in 1999, has held various business and editorial roles within Penton Media, and is currently editorial and custom strategy director of *Windows IT Pro, SQL Server Magazine,* and *System iNEWS*.

Spiceworks Fan
Sysprep Reminder
Group Policy Pointers
User Account Advice

LETTERS@WINDOWSITPRO.COM

[Editor's Note: This month's IT Community Forum presents a selection of comments and discussions from our website, duplicated here with only minor editing for style. You can take part in our community by logging on to www.windowsitpro.com.]

Spiceworks Fan

[The following comment is in response to Michael Dragone's "Spiceworks 4.5" review (InstantDoc ID 125235)]

[The current version of Spiceworks is 4.7], which, with the use of a script available from the Spiceworks website, [lets you] monitor both Exchange 2007 and Exchange 2010 servers. As a Spiceworks user since 4.1, it is a great tool, worthy of 5 out of 5 once you start using it on a proper network. I have it running on two subnets at the moment, one set as a remote collector to reduce traffic between the subnets. It may be a little tricky to get set up initially, but then again, so are other network management tools, a lot of which require the installation of agents on the clients. Spiceworks has zero footprint on the clients. Your login is also used to access the Spiceworks community, [which] has members from all over the world, so you could think of it as a free 24/7 support desk. In the four years that Spiceworks has been in existence, there are now over 1 million members/users. I am proud to be one of them.

—Lyons

Sysprep Advice

[The following comment is in response to Mark Minasi's Windows Power Tools column, "Creating Bootable VHDs with Disk2VHD" (InstantDoc ID 125422)]

Mark Minasi gives us excellent advice when he instructs us to generalize the images with Sysprep. We would not want to deploy the same non-Syspreped VHD to more than one machine. For example, the computer name should be changed. But Mark writes, "Creating an image and handing identical copies of that image out to zillions of machines can cause some trouble security-wise." Trouble security-wise? Hardly any. Please refer to "The Machine SID Duplication Myth" (blogs.technet.com/b/markrussinovich/archive/2009/11/03/3291024.aspx). I quote: "This blog post debunks the myth with facts by first describing the machine SID, explaining how Windows uses SIDs, and then showing that—with one exception—Windows never exposes a machine SID outside its computer, proving that it's okay to have systems with the same machine SID. Note that Sysprep resets other machine-specific state that, if duplicated, can cause problems for certain applications like Windows Server Update Services (WSUS), so Microsoft's support policy will still require cloned systems to be made unique with Sysprep." Bottom line: We should Sysprep, but we should not view it as a security measure.

—Dimitrios Kalemis

Group Policy Pointers

[The following discussion is in response to Harry Verge's "How to Trigger a One-Time Group Policy Registry Refresh" (InstantDoc ID 125265)]

I don't typically recommend [that folks tinker] with the registry metadata associated with Group Policy, simply because it's not a documented, supported method. But, that aside, the author misses some key points about how Group Policy processing works. First, the Group Policy engine has no notion of CSE versioning. This means that you really can't effectively isolate a refresh based on a particular CSE. For example, if you have a GPO (or several hundred GPOs) that contain both registry and security policy, if the Group Policy engine detects that it has to do work for either extension, then both will process, simply because it has no notion of which changed. So, you can't simply tweak the registry extension and assume that only that extension will process during the next refresh. Second, the notion of a toggle GPO is redundant. If you really want to trigger a refresh without touching every client, simply make a benign change to an existing GPO and then undo it (or use my GP-touch PowerShell cmdlet to increment a GPO's version number). Finally, it's a common misconception that a foreground refresh (i.e., reboot or re-logon) forces a Group Policy refresh. It does not. Group Policy will not refresh unless something *has changed—period. It doesn't matter if it's a reboot or a background refresh.*

—Darren Mar-Elia

I think this article is excellent. I also think that Darren Mar-Elia's criticism may be unfounded. First of all, the author did not

Detecting a Heartbeat

[The following comment is in response to Michael Otey's "Hyper-V Live Migration: A Step-by-Step Guide" (InstantDoc ID 125262)]

If the servers will be clustered, you should have a private network for the cluster heartbeat. This will either be a crossover cable between two hosts in a two-host cluster or a private VLAN on the switch dedicated just to these servers and this task. [According to Aidan Finn (www.aidanfinn.com/?p=10311)], "There are quotes out there about Windows Server 2008 failover clusters not needing a heartbeat network. But If CSV is configured, all cluster nodes must reside on the same non-routable network. CSV (specifically for re-directed I/O) is not supported if cluster nodes reside on separate, routed networks." Good article.

—Bolton

Windows IT Pro welcomes feedback about the magazine. Send comments to letters@windowsitpro.com, and include your full name, email address, and daytime phone number. We edit all letters and replies for style, length, and clarity.

really touch the subject of CSE versioning. Second, the notion of a toggle GPO is not redundant. Perhaps someone may not want to make a change to an existing GPO, even if that change would be a "benign" one. Third, I did not find anywhere in the article that the author supported the incorrect idea that a foreground refresh (i.e., reboot or re-logon) forces a Group Policy refresh.

—Dimitrios Kalemis

[My comment] was less criticism than correction. The lack of coverage of CSE versioning was precisely my point: Without understanding that, you risk unintended consequences by thinking you are just refreshing one CSE. The point about the "toggle GPO" is a matter of perspective, I suppose. I find it more invasive to create a new GPO than to "touch" an existing one. This becomes more poignant if you start talking about CSEs other than registry. There are many ways to force the Group Policy infrastructure to think a change has occurred. In my experience, the client-based ways don't necessarily scale. The point about reboot not necessarily triggering a refresh was probably me reading into what was written here: "After you run it, the registry client-side extension will refresh the registry at the next reboot or the next Group Policy refresh." Finally, while I don't necessarily see any problem with sharing this kind of information, it provides an incomplete picture of how Group Policy processing works under the covers. [As with] the CSE versioning point, unless you understand the full picture, you have the potential to cause some grief in your environment.

—Darren Mar-Elia

Running as Administrator

[The following discussion is in response to Russell Smith's "Solving User Account Problems" (InstantDoc ID 125157)]

Great article, Russ. We are trying to edit an executable [so that it doesn't] pop up the UAC prompt as part of [its] running in the login script. I tried using [Heaventools] to edit the manifest and change the <requestedExecutionLevel level="requireAdministrator" to <requestedExecutionLevel level="asInvoker" but now the file does not run automatically. It will run if I [right-click] and run as administrator but still pops up with the UAC prompt. Is there more to this I am missing?

—Anonymous

Thanks for the feedback. Maybe the executable really does require administrative privileges to run? So when you force it to start as a standard user, it fails to launch or it hangs. You might try using Process Monitor from the Sysinternals tool set to investigate the problem further.

—Russell Smith

Thanks, Russell, good article. I still have one situation I can't get around, [and] I was hoping you may have some ideas to help. When some machines are logged on with admin accounts I want some commands to run that are in the startup folder but they need elevated admin rights; instead of even prompting, they just fail to run, putting out a message to [the] command window: *The requested operation requires elevation (Run as administrator)*. I don't think I can change the manifest for the command prompt, if one even exists. For example, one command is a Netsh run from a .cmd file. I tried creating a shortcut to it and then under Properties, Compatibility tab, set the option *Run this program as administrator*, but this is always grayed out. Do you know of a way around this sort of issue without turning off UAC?

—Stenberg

The easiest way around this would be to set up a Group Policy startup script to run the Netsh command. GPO startup scripts run in the context of the local machine account, which has administrative privileges. As the name suggests, however, startup scripts run once before users log on. If the Netsh command must run as users log on, the Script Elevation PowerToy (technet.microsoft.com/en-us/magazine/2007.06.utilityspotlight.aspx) should do the trick.

—Russell Smith

InstantDoc ID 125837

An IT Pro Haiku

Laptop won't boot docked!
Hey, who set USB drives
First in the boot list?

—devinganger

NEED TO SECURE YOUR JOB, WANT TO START A NEW CAREER... WITHOUT BREAKING YOUR PIGGY BANK?

YOU CAN WITH TRAIN SIGNAL TRAINING COURSES FOR UNDER $400!

Our Computer Training Software:

- Is Scenario Based! It Mirrors Real World Challenges and Lays a Solid Foundation for Your Career!
- Helps You Prepare for Your Certification Exam the Right Way!
- 90-Day Money Back Guarantee!

We offer training in:

Windows, Cisco, VMware, CompTIA, Microsoft, CWNA

Get a headstart and call our toll free number today!

(888) 229-5055

TRAINSIGNAL

THE GLOBAL LEADER IN PROFESSIONAL COMPUTER TRAINING™

www.trainsignal.com

Thurrott

"Windows Phone, unlike Google Android and Apple iPhone, doesn't support much in the way of PC-based sync. Microsoft is providing no interfaces for synchronizing with Outlook."

What You Need to Know About Small Business Server "Aurora," the Windows Phone 7 Debate, IE 9, and More

A year after Windows 7's release, Microsoft is starting to update its other product lines, including two new Small Business Server (SBS) offerings, a Windows Home Server (WHS) upgrade, Windows Phone 7, and Internet Explorer (IE) 9. Let's examine each of these upcoming releases, as well as Windows 7's deployment record with businesses, and look at the return of the Slate PC. Here's what you need to know.

Small Business Server "Aurora" Update

As Microsoft ships a near-final, public release of its upcoming Windows Small Business Server product, code-named Aurora, in late August, it appears that this intriguing "cross-premises" server offering will ship in final form by the end of the year. With the caveat that I sometimes feel like the patron saint of small business computing, I think Aurora could be a game changer: This amazing product offers all of the Active Directory (AD)-based identity, security, and computer management that businesses need but without the complexity. And, I expect, without the cost: Though Microsoft hasn't yet revealed Aurora pricing, it's going to have to come in well under regular SBS pricing. This could be a mainstream, high volume product.

Of course, the success of Aurora won't hinge on pricing alone. What makes this product really work is its logical targeting of the real needs of small businesses. There's no need for a true admin or an IT pro; instead, the simplified system administration can be easily delegated to different people in the office. New users can add their own PCs to the domain, and Aurora automatically copies all their data and settings over to a new domain account. And Aurora, like the WHS products on which it's based, offers excellent, centralized image and file backup of all connected PCs.

Where Aurora really shines, however, is storage: Using WHS's Drive Extender solution, newly added internal and external hard drives can be added to a bottomless pool of storage that doesn't need drive letters and offers data duplication functionality at the share level. Everything a small business really needs is available on site.

The "cross-premises" promise of Aurora means that additional services—email, calendaring, communications, document collaboration, and more—can be added either via traditional on-premises servers (Exchange, SQL Server, and others) or via cloud-based services such as Exchange Online and SharePoint Online. And, of course, you can mix and match as your needs—and checkbook—allow. After the misguided disaster of Windows Essential Business Server, Aurora is a breath of fresh air and proof that Microsoft really does understand its different business market segments.

WHS "Vail" Update

In tandem with the near-final Aurora code release, Microsoft also issued a second preview release of its upcoming WHS version, code-named Vail. This product looks and works much like Aurora, but with two important distinctions: It utilizes workgroups and Windows 7-based homegroups instead of a true domain, and its storage features are oriented around media sharing, as you might expect of a consumer solution.

But don't write off Vail so quickly: In many ways, Vail is an ideal solution for very small businesses as well as individuals, and if even vastly simplified domain management seems like overkill, this could be an interesting solution. But the real advantage of Vail over Aurora is interoperability: An Aurora server must be the first server in a new domain, so you can't add Aurora to existing domain environments. But if all you're looking for is the amazing storage functionality in Aurora, Vail features Drive Extender, too. So it's an ideal storage solution for any smaller environment.

The Great Windows Phone 7 Debate

Microsoft and its partners will launch Windows Phone 7 this October, and although we've discussed the software giant's innovative new smartphone platform here in the past, there's an interesting debate emerging in the days leading up to the launch. Windows Phone, unlike its predecessor, Windows Mobile, and unlike leading competitors such as Google Android and Apple iPhone, doesn't support much in the way of PC-based sync. That is, while you can—in fact, must—use the Zune PC software to synchronize media with the devices, Microsoft is providing no interfaces for synchronizing with Outlook or other desktop-based productivity solutions. And this is causing some predictably painful reactions in certain quarters.

The solutions Microsoft does support are cloud-based in the sense that they connect with Windows Phone over the air. These include Exchange and Exchange Online, Gmail/Google Calendar, Windows Live/Hotmail, and to a lesser extent Facebook and Yahoo! Mail. Microsoft is also supporting any IMAP- or POP3-based email accounts as well. By supporting only direct connectivity between the phone and online accounts, Microsoft is shutting out PC-based middlemen, and not just Outlook, but also previous generation sync solutions like ActiveSync or Windows Mobile Device Center.

This is a bit forward-looking. But it's also the right decision, arguably, given recent industry trends and the dangers associated with relying on a single PC as your central data store.

Windows Phone is going to be less than ideal for those businesses that require the full set of AD and Group Policy device management features. That's because unlike Windows Mobile, Windows Phone will support only a subset of those features at launch, meaning that Microsoft's previous generation mobile platform will stick around for an additional year or two.

Windows 7 Racks Up Impressive Numbers

The more time that goes by, the better Windows 7 looks. By the end of July, Microsoft had sold over 175 million copies of the OS, at a rate approaching 10 units per second. Amazingly, sales are actually accelerating, in part because of generally improved PC sales. Microsoft CEO Steve Ballmer said recently that PC makers will sell over 400 million PCs in calendar year 2011, up from the 360-370 million previously predicted.

But the best news for Microsoft, perhaps, is that businesses are biting: Over 65 percent of enterprises are already migrating to Windows 7, and many of them are jumping on board with Office 2010 as well. In an admittedly unscientific poll of my own readers, I was swamped by positive stories about Windows 7 migration experiences, with only a few reporting they had to hold off because of compatibility reasons with apps used by specific user groups. There are lots of factors at play here, of course, including an improving economy, a near-decade's reliance on an increasingly creaky Windows XP, and a general consensus that Windows 7 really is better than previous Windows versions, with important productivity gains. But regardless of the reasons, it's pretty clear Microsoft is looking at its biggest success story in a long, long time.

Slate PC Returns from the Dead

With Apple's iPad selling well despite some frustrating problems, Microsoft and its PC maker partners are racing to provide suitable alternatives. Well, maybe "racing" isn't the right word. We're looking at early 2011 before the real iPad killers arrive. And by then it might be too late.

So why the 2011 requirement? That's when Intel will release its Oak Trail "system on a chip," which will apparently offer much better power management and performance than the company's current Atom and i-series chips. This is amazingly bad timing for Microsoft, as all of its partners appeared to be caught by surprise by the iPad despite months of warning.

Since we're looking at 2011 anyway, Microsoft should forget about Windows on a Slate PC—that ship sailed several years ago, when virtually no one bought any of the first generation Tablet PCs—and think about porting its well-executed Windows Phone 7 OS to a tablet-type device. Unlike Apple's iOS behind the iPad, Windows Phone OS would actually make sense on a tablet, and, in fact, seems like it was designed from the get-go for that kind of form factor. Sadly, Microsoft continues to insist it has no plans for such a thing.

IE Might Actually Be a Contender

Maybe I'm preaching to the choir here—after all, IE remains the de facto choice in corporate environments—but it appears as if Microsoft's much-maligned web browser might be poised for a comeback. Usage in IE had been on a near-linear freefall for a few years, thanks first to Mozilla Firefox and then, more recently, because of a surge of interest in Google's Chrome browser.

But that all changed in 2010. First, IE's overall market share decline evened out, then stopped falling altogether, and since May 2010, the browser has gained market share, while both Firefox and Chrome have stalled. Meanwhile, IE 8 has outgrown the competition, and is the most frequently used web browser.

Critics will point out that IE 8 is bundled with Windows 7, and that Microsoft's OS popularity must be helping the browser. Fair enough, but IE 8 growth has also come during a period in which European Union-based Windows installations have triggered a browser ballot screen any time a user picks IE as the default browser. Combined with the heightened competition, then, IE's recent successes have come during a time when the browser has faced unparalleled opposition.

And now IE 9 is on the way. Through much of 2010, IE 9 was almost a science experiment, with Microsoft focusing on the developer-oriented underpinnings of the product. The effort paid off: Not only has Microsoft embraced the standards-based web, it's also found a way to leverage unique Windows integration features like hardware acceleration to make its next browser even more desirable.

By the time you read this, Microsoft should have debuted the first IE 9 public beta. My sources tell me that IE 9 will be lean and mean in ways that the bloated IE 8 isn't, and I hope that's true: One of the selling points of Google Chrome, in particular, is its stripped-down UI that gets out of the way and lets the content you're viewing take center stage. If Microsoft can duplicate that experience and combine it with its standards effort and hardware-accelerated rendering, IE 9 could be a much bigger deal than previously expected. And that's an amazing turnaround for a product that many of us had written off less than a year ago.

PDC 2010

Microsoft sporadically holds a developer-oriented event called the Professional Developers Conference (PDC). This year, it's on October 28 and 29, in Redmond for the first time. It will focus on cloud services, mobile development, tools, browser strategy, and gaming. I'll be going, so if you're in Redmond then, drop me a line (thurrott@gmail.com).

InstantDoc ID 125858

PAUL THURROTT (thurrott@windowsitpro.com) is the news editor for *Windows IT Pro*. He writes a weekly editorial for *Windows IT Pro UPDATE* (www.windowsitpro.com/email) and a daily Windows news and information newsletter called *WinInfo Daily UPDATE* (www.wininformant.com).

Lotus knows you can't truly appreciate a cloud from behind windows.

Not all clouds are equal. Unlike some vendors, LotusLive™ makes it easy to collaborate with people outside your company. Plus, you get a comprehensive suite of integrated productivity tools, including e-mail, file sharing, social networking, web conferencing, instant messaging and integrated third-party apps. All of this for only $10 per user per month. Does your cloud offer all that?

Smarter software for a Smarter Planet.

lotusknows.com

Minasi

"Since the mid-1990s, IT pros needing to perform quick OS installs/reinstalls have relied on *disk imaging*."

ImageX Provides Disk Imaging on a Budget

Use this free Microsoft command-line tool to quickly image new or troubled computers

In "Creating Bootable VHDs with Disk2VHD" (InstantDoc ID 125422), I talked about how to boot a system from a VHD file and mentioned that you need some kind of tool to create an image from which you can make that VHD. I gave you an overview of Sysinternals' Disk2VHD tool and suggested that other tools were available, as well. This month, let's take a look at a personal favorite: ImageX.

Since the mid-1990s, IT pros needing to perform quick OS installs/reinstalls have relied on *disk imaging*, which lets you configure a "reference" Windows system just the way you want it, then preserve all of that hard disk's files, folders, and metadata into one large file: the *disk image* file. You can then use a tool to quickly copy that image onto the troubled PC and restore it to service. The premier tool in that arena has been the popular Symantec Ghost. (See this month's comparative review, "3 Disk Imaging Products—Redux," InstantDoc ID 125797, for a look at Ghost, as well as Acronis Snap Deploy and Paragon Deployment Manager.) A few years ago, however, Microsoft began giving away a command-line replacement for Ghost called ImageX, which is part of Microsoft's Windows Automated Installation Kit (WAIK).

Suppose you've spent a week or two creating the perfect Windows 7 system, including service packs, hotfixes, applications, and application settings, and you want to make dozens, hundreds, or thousands of PCs run just like that system. First, you've got to *generalize* that system with Sysprep, as I discussed in "Get Ready for Imaging with Sysprep" (InstantDoc ID 125532). After that, your reference system's hard disk is *almost* ready to be imaged. Most imaging tools can't image an OS while that OS is running, so you'll need to boot the reference system from some *other* OS before you can fire up ImageX. Now, you could do that by putting a second copy of Windows on the reference image and booting from that, but that's sort of cumbersome. Microsoft's answer is the simple, streamlined Windows Pre-installation Environment (WinPE). You get WinPE in the WAIK, so your next task is to set up a bootable copy of WinPE on either a CD or USB drive. For detailed instructions, see the Microsoft article "Walkthrough: Create a Custom Windows PE Image" (technet.microsoft.com/en-us/library/dd744533%28WS.10%29.aspx).

Once booted to WinPE, you'll need access to the imagex .exe program file because the base WinPE image doesn't include imagex.exe. Probably the easiest way to get imagex.exe onto your WinPE system is to grab it from a WAIK-enabled system, then copy it to X:\Windows on your WinPE system. (WinPE doesn't store any changes from boot to boot, so you'll need to re-copy that imagex.exe file to X:\Windows every time you need to do some imaging.)

Because disk image files are pretty big, you'll also need a place to copy the imaged hard disk to, although if you have a lot of free space on your C drive, ImageX—unlike most disk-imaging tools I've worked with—will let you image the C drive onto itself. For the purpose of this article, let's assume that I've booted into WinPE on my reference system and then either mapped to a network share or connected an external USB drive to my reference system and called that drive G.

You'll find when you boot your system under WinPE that you're running from drive X. Your goal is to capture an image of drive C from the reference system onto G, resulting in a file on G that you can deploy to other systems. Now, when booted from WinPE, your system might rearrange its drive letters, so what *was* drive C on your reference system when it was booted in Windows 7 might have changed under WinPE. So, take a moment to ascertain which drive letter you want to image. (I'll assume that it's drive C.)

Now you're ready to capture a disk image. The basic ImageX command to create a new image is

```
imagex /capture <sourcedrive> <targetdrive:\filename>.wim
  "<description>"
```

For example, to capture the C drive onto G, creating a file named baseimage.wim, you'd type

```
imagex /capture c: g:\baseimage.wim "Base Win 7 image for
  our organization"
```

The result will be a file named G:\baseimage.wim. The file will be big but not as large as your initial C drive—for a couple reasons. First, ImageX doesn't bother copying your page file, which as you know can be quite large, and second, ImageX by default applies a little compression to the image. As you'll see next month, you can control the level of ImageX's compression and a whole lot more with ImageX's many options.

InstantDoc ID 125743

MARK MINASI (www.minasi.com/gethelp) is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 25 books, including *Administering Windows Vista Security: The Big Surprises* (Sybex). He writes and speaks around the world about Windows networking.

VISION SOLUTIONS WELCOMES DOUBLE-TAKE SOFTWARE

Double-Take ⇕ iTERA ⇕ MIMIX

Acquisition Creates an Information Availability Powerhouse for Windows, Linux, Power Systems and Cloud Computing

Vision Solutions now offers a best-in-class portfolio of high availability and disaster recovery software for Windows®, Linux®, Power Systems™ and Cloud Computing.

The New Vision Solutions features a greater choice of technologies and one of the world's most innovative R&D, service and global customer support teams for availability software.

Take a closer look at visionsolutions.com or call 800-957-4511.

Vision CEO Welcomes Double-Take

Visit visionsolutions.com

VISION® SOLUTIONS

Easy. Affordable. Innovative. *Leaders Have Vision.*

Otey

"One of the most common yet hard-to-find things that you'll need to do in a remote desktop session is to send a Ctrl+Alt+Del signal to the remote system."

Remote Desktop Keyboard Shortcuts

Let your fingers take screenshots, switch programs, and other useful admin tasks—remotely

Have you ever been using one of your RDP sessions to manage a remote system and found that you needed to press Ctrl+Alt+Del on the remote system but you couldn't remember how to do it inside the Remote Desktop window? Just pressing Ctrl+Alt+Del sends the keystrokes to your local desktop, which certainly isn't what you wanted. If you've ever had this problem, then this Top 10 column is your answer. In this column, I show you ten handy keyboard shortcuts you can use in your remote desktop sessions.

10 **Ctrl+Alt+plus sign (+)—**Dealing with capturing screen images from a Remote Desktop session can be a mystery. If you press Print Screen, you get an image of your local desktop—not the remote desktop. Pressing the Ctrl+Alt+plus sign (+) keyboard shortcut captures a snapshot of the entire client window area of Remote Desktop and is the same as pressing Print Screen on your local desktop.

9 **Ctrl+Alt+minus sign (-)—**Sometimes you don't want an image of the entire desktop; sometimes you want just a selected window. Pressing the Ctrl+Alt+minus sign (-) keyboard shortcut captures a snapshot of just the active window within the remote desktop session. This key combination is the same as pressing Alt+Print Screen on your local desktop.

8 **Alt+Home—**Pressing the Alt+Home keyboard combination with Remote Desktop displays the Start menu on the remote system. The Start menu gives you quick access to the different programs installed on the remote system. This key combination is the same as pressing the Windows key on your local desktop.

7 **Alt+Delete—**Pressing the Alt+Delete keyboard combination in the remote desktop session opens the Windows menu of an application running on the remote system. The Windows menu is typically displayed under the icon in the extreme upper left corner of most Windows applications, and it lets you move and resize the application.

6 **Ctrl+Alt+Break—**Sometimes you might want the Remote Desktop window to be displayed in full-screen mode just as if you were using your local desktop. If you want to toggle the remote desktop session between a window and a full-screen display, you can press the Ctrl+Alt+Break keyboard combination.

5 **Ctrl+Alt+Pause—**Like the previous item, the Ctrl+Alt+Pause keyboard combination switches between full screen and windowed mode. However, with this keyboard shortcut, the remote desktop window remains at its standard size and doesn't fill the entire local desktop. Instead, it's displayed on a black background.

4 **Alt+Insert—**Sometimes you want a quick way to switch between the different programs that you have running. Pressing the Alt+Insert keyboard combination lets you cycle through the programs on the remote system in the order that they were opened. This process is the same as using Alt+Tab on your local desktop.

3 **Alt+Page Down—**Another way to cycle through the running programs on your Remote Desktop session is to use the Alt+Page Down keyboard shortcut. Pressing this key combination lets you switch between programs on the remote desktop session, moving from right to left in the Windows task switcher. This is the same as Alt+Shift+Tab on your standard desktop.

2 **Alt+Page Up—**Pressing Alt+Page Up lets you switch between programs on the Remote Desktop session, moving from left to right in the Windows task switcher. This is the same as Alt+Tab on your standard desktop.

1 **Ctrl+Alt+End—**One of the most common yet hard-to-find things that you'll need to do in a remote desktop session is to send a Ctrl+Alt+Del signal to the remote system. Press Ctrl+Alt+End if you need to send a Ctrl+Alt+Del keystroke combination to the remote system. This keystroke opens the Microsoft Windows Security dialog box, which lets you lock the computer, log off, change your password, and start Task Manager.

InstantDoc ID 125784

MICHAEL OTEY (motey@windowsitpro.com) is technical director for *Windows IT Pro* and *SQL Server Magazine* and author of *Microsoft SQL Server 2008 New Features* (Osborne/McGraw-Hill).

Wheeler

"I'll show you how to view incoming WMI queries as close to real time as possible, then correlate where the query is coming from."

An Easier Way to View Incoming WMI Queries

A helpful script outputs timestamped WMI queries in an easy-to-read format

Windows Management Instrumentation (WMI) is a widely used technology on both Windows server and client, employed for tasks such as inventory of hardware and software, determining whether services or processes are running, and many others. Usually, WMI just works, but when it doesn't and you need to investigate it, doing so can be very difficult. For instance, one of the most common problems with WMI is high CPU spikes for the one of the provider processes (wmiprvse.exe). The provider process is where the WMI queries are executed, and the WMI service process (svchost.exe) is where the query results are returned to the process that executed the WMI query. As an administrator, you've probably either run into a high-CPU–usage issue with WMI or needed to view incoming WMI queries. In the course of investigating a high-CPU issue, it would be very beneficial to see the queries being executed and determine where the query came from.

Unfortunately, viewing incoming WMI queries is at best a cumbersome process that can frustrate administrators trying to answer the simple question "What WMI queries are being executed on my system?" In this article, I'll show you how to view the incoming queries as close to real time as possible, then correlate where the query is coming from. This will enable you to immediately troubleshoot common WMI high-CPU issues without having to perform complicated debugging steps or scan through hundreds of log file events—or make that dreaded support call.

Retrieving WMI Trace Logs

As described in the blog post "WMI Debug Logging," tinyurl.com/3a6b7ls, you can turn on the tracing mechanism inside the Event Viewer to view the tracing logs for WMI. However, because of the large number of events collected, it's difficult to determine the latest incoming queries, when they executed, and where they are coming from. Even the filtering mechanism built into the Event Viewer menu does not provide the ability to filter only on "WMI queries".

Enter wevtutil.exe, the Windows events command-line utility. Wevtutil, included in Windows Vista and later, lets you retrieve information about event logs as well as export, run queries, and set properties on the log files. For instance, Wevtutil can set the tracing flag on the WMI-Activity event log either locally or remotely, so that you can start receiving WMI trace logs in the Event Viewer interface.

Here is the exact syntax to set the tracing flag on the WMI-Activity log file (start an elevated command prompt), along with a message displayed after you enter the command:

```
C:\Windows\system32>wevtutil sl Microsoft-Windows-WMI-
  Activity/Trace /e:true

**** Warning: Enabling this type of log clears it.  Do you
  want to enable and
clear this log? [y/n]:
y
```

Notice that you are asked to clear the log, which is required to enable the tracing option. (Note: It is not required to clear the log to disable the tracing flag.) There is also an /r option that lets you set the tracing flag on a remote system.

After the tracing flag is set, you will start to receive "Informational" events in the WMI-Activity event log. However, the verbosity is high, and even if you wanted to use the filtering mechanism, you still would not be able to filter for just WMI queries. Again we will use wevtutil.exe to help us manually view the incoming WMI queries. Here is the syntax to view the WMI queries using Wevtutil (after enabling the tracing flag and from an elevated command prompt).

```
<Event xmlns='http://schemas.microsoft.com/win/2004/08/
  events/event'><System><Provider Name='Microsoft-
  Windows-WMI-Activity' Guid='{1418ef04-b0b4-4623-bf7e-
  d74ab47bbdaa}'/><EventID>1</EventID><Version>0
  </Version><Level>4</Level><Task>0</Task><Opcode>0</Opco
  de><Keywords>0x8000000000000000</Keywords><TimeCreated
  SystemTime='2010-08-20T17:26:52.9118705002'/><Event
  RecordID>17</EventRecordID><Correlation/><Executio
  n ProcessID='1156' ThreadID='8136' ProcessorID='1'
  KernelTime='25' UserTime='28'/><Channel>Microsoft-Windows-
  WMI-Activity/Trace</Channel><Computer>Computer1
  .na.corp.contoso.com</Computer><Security UserID='S-
  1-5-18'/></System><UserData><Operation_ xmlns:auto-
  ns2='http://schemas.microsoft.com/win/2004/08/
  events' xmlns='http://manifests.contoso.com/win/2006/
  windows/WMI'><GroupOperationId>467724</GroupOperation
  Id><OperationId>467725</OperationId><Operation>Start
  IWbemServices::ExecQuery - select * from Win32_
  OperatingSystem</Operation><ClientMachine>Computer1
  </ClientMachine><User>na/user1</User><ClientProcessId>5900
  </ClientProcessId><NamespaceName>\\.\root\cimv2</
  NamespaceName></Operation_></UserData></Event>
```

Figure 1: Wevtutil output

```
GroupOperationId = 467724; OperationId = 467725; Operation = Start
  IWbemServices::ExecQuery - select * from Win32_OperatingSystem; ClientMachine =
  User1; User = NA\user1; ClientProcessId = 5900; NamespaceName = \\.\root\cimv2
```

Figure 2: Wevtutil output in text format

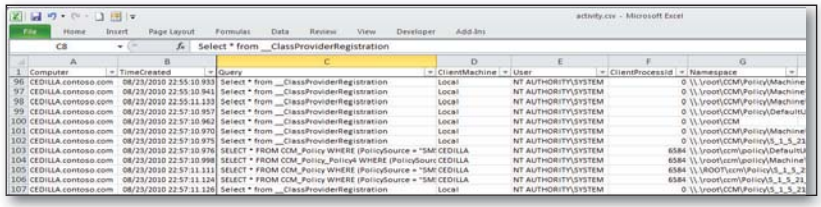

Figure 3: Output from WMIActivity.vbs script

```
c:\>wevtutil qe Microsoft-Windows-WMI-
  Activity/Trace /q:"*[System[(Level=4
  or Level=0) and (EventID=1)]]"
  |findstr /i execquery
```

```
C:\>wevtutil qe Microsoft-Windows-WMI-
  Activity/Trace /q:"*[System[(Level=4
  or Level=0) and (EventID=1)]]"
  /f:text |findstr /i execquery
```

The first option qe means "query events from a log file"; then the name of the log file is entered (Microsoft-Windows-WMI-Activity/Trace). The /q: option says to filter the event log for just those specific events; in this case we are asking for Level 4 or Level 0 AND Event ID=1, which is the event ID that all WMI queries are logged as. Then we pipe the results to findstr (find string) and filter for just the /execquery commands.

The output from the previous command looks similar to that shown in Figure 1. This output is for just one query—imagine if you received hundreds of queries! That would be a lot of data to sort through. However, we can change the format of this output from XML, to text by making a small change to the Wevtutil command:

Notice the change we made was simply adding the /f: option, which means "format," and supplying the format as text. Now our output, shown in Figure 2, is more manageable. Although this output is less verbose, it still too verbose for the simple task of viewing incoming WMI queries, plus it doesn't provide any timestamp to indicate when the query occurred.

Automating WMI Activity Tracking

It would be helpful to have a script that you can schedule to run and would provide manageable output, show timestamps to indicate when a query occurred, and let you see the latest queries since the last time the script was run. And that is exactly what the Microsoft support team has done for you. We've created a script that shows you the incoming WMI queries and when they occurred. Depending on the version of Windows used, you'll be able to tell which machine and process ID the query came from, either locally or remotely.

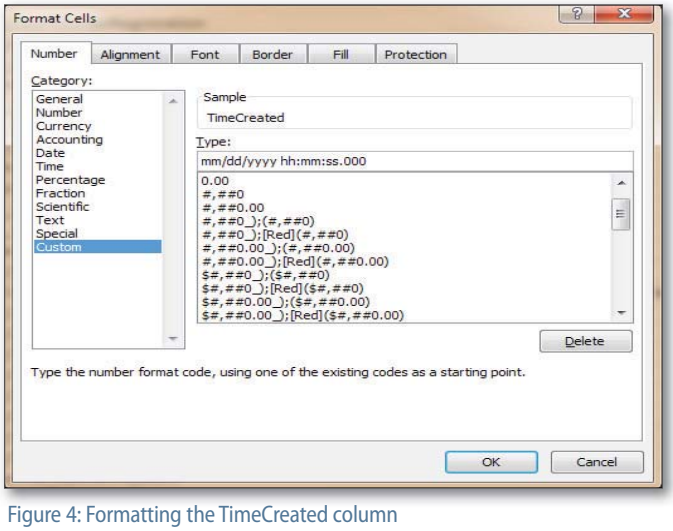

Figure 4: Formatting the TimeCreated column

The script, WMIActivity.vbs, queries for WMI events and outputs only the WMI query activity. (You can download the script at www.windowsitpro.com, InstantDoc ID 125876; simply rename the text file containing the script with a .vbs extension to make it executable.)

Figure 3 shows the output in Excel spreadsheet format. Notice that the output includes several columns of data that give you as the administrator insight into when the query executed, the query itself, the machine where the query came from (ClientMachine), and the Client Process ID.

I should mention that the ClientMachine and ClientProcessID columns will be populated only for systems running Windows 7 or Windows Server 2008 R2 or later. Vista and Windows Server 2008 don't include the code to capture the client that sent the query or the client process ID.

The first time the script is run, a timestamp is put into the registry under the key HKLM\Software\WinITPro. Each time the script is run, the timestamp is checked and only the log entries newer than the timestamp are pulled. That way you see only the newest queries rather than having to wade through much of the same information or navigate through a cumbersome Event Viewer–style menu.

You will need to format the TimeCreated column with a custom format. You can do so by selecting the entire column, right-clicking, and selecting Format Cell, as Figure 4 shows. You could also create a pivot table, so that you can view the queries from a "Count" perspective.

A Little Extra Troubleshooting Help

With help from the WMIActivity.vbs script file, you'll be able to view the WMI queries being sent to your system as close to real time as possible. This script will give you some extra help in troubleshooting WMI high-CPU issues as well as better insight into what and who is interrogating your system.

InstantDoc ID 125876

SEAN WHEELER (seanw@microsoft.com) is a senior premier field engineer in Microsoft's Premier Support team, assigned to support some of the largest enterprise customers. He is one of the original creators of the MPS Reporting tool. He specializes in scripting, debugging, and performance-related issues.

Identifying Logons from Remote Machines

READER TO READER

Batch File Finds Out Who RDPed the Server

Administrators and privileged users often use RDP client software such as Remote Desktop to access a server to perform tasks and troubleshoot problems. When doing so, it's difficult to pinpoint the last actions that were applied to the server. However, if you know which administrators and privileged users recently logged on to that server, you can contact them to find out this information. I've come up with a batch file—RDPLogon.bat—that lets you quickly identify those people.

Here's what you need to do:

1. On your server, create a folder named RDP on the C drive (C:\RDP). For the batch file to work, the Remote Desktop Users group must have permission to write to this location. (It's assumed that the people who use RDP to access servers from remote machines are either administrators—who by default have RDP permission—or members of the Remote Desktop Users group.)

2. Create a batch file that contains the code shown in Listing 1. Alternatively, you can download RDPLogon.bat from the *Windows IT Pro* website by going to www.windowsitpro.com, entering 125864 in the InstantDoc ID box, clicking Go, then clicking the *Download the Code Here* button. Note that your server needs to be using the default RDP port (port 3389) for RDPLogon.bat to work.

3. Place the batch file in the C:\Documents and Settings\All Users\Start Menu\Programs\Startup folder on the server.

James Lim

RDPLogon.bat captures information about each person who logs on to the server and appends that information to an output file named RDPlog.txt, which resides in C:\RDP. By looking at the RDPlog.txt file, you can identify the administrators and privileged users who logged on recently and which machines they used for that logon.

Figure 1 shows sample output from RDPLogon.bat. This output contains information about three logons. (For easier reading, the batch file puts a dashed line after each logon entry.) The first line in each entry shows when the logon occurred (date and time) and by whom (user's domain ID).

You can find out where each logon occurred by looking at the second line in each entry. The information in this line is obtained using the Netstat command. The second string of numbers, which I highlighted in yellow in Figure 1, shows the IP address and RDP port number of the remote machine used for the logon. (The first string of numbers is the server's IP address and RDP port number.) If an entry is missing the Netstat results, like the last entry in Figure 1, it means that the logon was done from the physical server console.

```
Fri 08/06/2010  9:13:38.87 Popeye.Sailor
  TCP    155.14.246.208:3389    155.14.230.71:2386     ESTABLISHED
============================================================

Fri 08/06/2010  9:20:18.96 Olive.Oyl
  TCP    155.14.246.208:3389    155.14.230.71:2398     ESTABLISHED
============================================================

Fri 08/06/2010  9:32:35.27 Brutus.Bully

============================================================
```

Figure 1: Sample RDPlog.txt file

RDPLogon.bat is a simple tool that you can use to quickly identify the administrators and privileged users who recently logged on to a server. Armed with this knowledge, you can contact them to find out the last actions they performed on the server. This is much faster than looking through and trying to decipher huge event logs to find the last actions. Note that this tool is meant for environments in which administrators and privileged users are cooperative, since they could conceivably delete their logon entries if they know the RDPlog.txt file exists and want to hide their actions.

—James Lim, systems manager, Distributed Systems and Services at Neptune Orient Lines

InstantDoc ID 125864

Listing 1: RDPLogon.bat

```
@Echo Off

Echo %date% %time% %username% >> c:\RDP\RDPlog.txt
Netstat -n -p tcp | Find ":3389" >> c:\RDP\RDPlog.txt

Echo============================================================================================================ >> c:\RDP\RDPlog.txt
```

Tell the IT community about the free tools you use, your solutions to problems, or the discoveries you've made. Email your contributions to r2r@windowsitpro.com.

If we print your submission, you'll get $100.

Submissions and listings are available online at www.windowsitpro.com. Enter the InstantDoc ID in the InstantDoc ID text box.

WINDOWS IT PRO VIP is

Educational—with FREE eLearning courses and eBooks available 24×7

Deep—housing over 41,000 articles on DVD and online, some exclusively for VIP members

Broad—solutions, tips, and tricks for any Windows or SQL Server issue that can stump you

In fact, Windows IT Pro VIP delivers more than **$1,000 of resources and expertise for just $199 a year.**

HOW WINDOWS IT PRO VIP BEATS A SEARCH ENGINE		
	Windows IT Pro VIP Delivers:	**Search Engines Deliver:**
Reliability	Road-tested advice from experts who put their reputation on the line	Well-meaning but potentially harmful tips in the latest Wikipedia entry
Speed	The answers you need in seconds searching by keyword, topic, or publication	Lost time spent perusing sites that have mastered search engine rankings but not the art of Active Directory or patch management
Impartiality	Authors and experts who challenge the Microsoft party line and influence industry change	Conventional wisdom touted by industry insiders afraid to tell it like it is

Order Online Now at windowsitpro.com/go/vip

- Domain Controllers
- BranchCache
- DNS
- Outlook
- ESX

ANSWERS TO YOUR QUESTIONS

Q: Is it a good idea to virtualize all of my domain controllers (DCs)?

A: Jay Hurt of Intrawest Colorado sent in this question:

"We're about to replace our two domain controllers and we want to run them as VMs. But we're getting a lot of resistance from our corporate IT team. They insist that we run the DCs on physical servers. Is this a good idea? Are there best practices associated with running domain controllers as virtual machines?"

In fact, there are. Plenty of companies today have moved portions of their Active Directory processing to virtual machines (VMs). Because DC processing tends to involve relatively light resource use in small-and medium-sized environments, DCs can be a great candidate for virtualization.

Jay tells me that his two DCs are two in a forest of many others. For environments of this size, I tend to advise people to virtualize, but keep an eye on your performance. Never snapshot a DC (never, never!). And always keep at least one physical DC per site.

Keeping at least one physical DC per site prevents you from getting into the circular problem where your virtual environment relies on your DC while your DC relies on your virtual environment. Should your virtual environment experience a problem, you'll still have that physical DC around to respond to requests.

Monitoring performance is also critical, though today's well-managed virtual hardware tends to give you enough horsepower for the job. While DC processing needs for smaller environments tend to be low, Exchange and other services will increase that load.

Check out the Microsoft article "Things to Consider When You Host Active Directory domain controllers in virtual hosting environments" (support.microsoft.com/kb/888794) for a few more thoughtful suggestions for virtualizing your DCs.

—Greg Shields

InstantDoc ID 125670

Q: What's a hashtable in PowerShell?

A: Also called a dictionary or associative array, a hashtable is simply a set of key/value pairs. If you know a key, you can look up its value. It's similar to using a dictionary—if you know a word (the key), you can look up its definition (the value).

To create a hashtable, you use the array operator, @, along with curly braces {}. Separate key/value pairs with a semicolon (;). Try running these lines in the shell:

```
$hash = @{'one'='don';'two'='greg'}
$hash.don
$hash.greg
$hash
```

Perhaps the most common use of a hashtable is with the Select-Object and Format-Table cmdlets, where they can be used to create custom columns in your output. Those cmdlets look for keys named "Name" ("Label" is an alternate; you can also use "n" or "l") and "Expression" (or "e"):

```
Get-WmiObject -class Win32_LogicalDisk |
Format-Table DeviceID,
  @{n="Size(GB)";e={$_.Size / 1GB -as
  [int]}}
```

—Don Jones

InstantDoc ID 125729

Q: I'm creating a server that will be sent to a branch location in my datacenter. It will be a hosted BranchCache server and I want to pre-cache some content. How do I do this?

A: There isn't currently a command you can run to pre-cache content onto a server. The best solution today would be:

1. Enable the hosted BranchCache feature on the server.
2. Configure one client to use that server for its hosted BranchCache. (You could also just run directly on the hosted cache server, and it will offer data to itself for caching.)
3. On the client, read all the data you want to be cached on the hosted BranchCache. You could script this; download the file, delete it straight away, then move on to the next file.
4. Check that the hosted BranchCache feature's cache is increasing in size, which means the data is being cached.

Obviously this isn't ideal, but it's a workable solution.

—John Savill

InstantDoc ID 125755





Jan De Clercq | jan.declercq@hp.com
Don Jones | powershell@concentratedtech.com
William Lefkovics | william@mojavemediagroup.com
John Savill | jsavill@windowsitpro.com
Greg Shields | virtualgreg@concentratedtech.com

Q: What are some shortcuts for locking the Windows desktop without using the Start menu?

A: Besides the Lock option on the Windows Start Menu, you can also simultaneously press the Windows and L keys on your keyboard. You can also use a desktop shortcut link. To create a desktop shortcut for locking your desktop, follow these steps:

1. Right-click your desktop and select New, Shortcut to start the Create Shortcut wizard.

On the first page of the wizard, enter the following path in the location of the item box:

```
rundll32.exe user32.dll,LockWorkStation
```

2. Click Next.
3. On the next page, type a name for the shortcut as you want it to appear on your desktop.
4. Click Finish.

To change the icon of your shortcut—the default icon is a file symbol—you can use these additional configuration steps:

1. Right-click the shortcut and select Properties.
2. On the Shortcut page, click the Change Icon… button.

On the Change Icon page, enter the following path:

```
%SystemRoot%\system32\shell32.dll
```

3. Select whatever icon you prefer (there several key and lock icons in the list) and click OK.
4. Click Apply.

Remember that in Windows Active Directory domain environments, administrators can also use Group Policy Object (GPO) settings to automatically lock user desktops after a certain amount of idle time and password protect them. To do so, they must use the Enable screen saver, Screen saver timeout and Password protect the screen saver settings in the User Configuration\Administrative Templates\Control Panel\Personalization GPO container.

—Jan De Clercq

InstantDoc ID 125768

Q: How can OneNote 2010 be integrated with Outlook 2010?

A: Microsoft OneNote is a virtual notebook binder for different types of content. You can use it to store text and content from almost any copyable source. OneNote provides a free-form intermediary between different applications and formats. It uses the familiar nomenclature of pages and notebooks for compartmentalizing information. Content in OneNote can also be shared for collaboration, saved in a network location, such as SharePoint, or on the web, in a location such as Office Live. It can be used as a rudimentary database of basic documents and an intermediary for distributing information, even in heterogeneous environments.

OneNote offers some basic free-form editing, highlighting, and other customizations. Content, including email messages, can be saved to OneNote, manipulated in ways unavailable within Outlook, and then emailed again. OneNote can save pages in many different formats, including Word .docx and .doc, .pdf, and .xps. OneNote can also save content as .xml or in its own .one format. The OneNote extension .one is the same for OneNote 2007 as it is for OneNote 2010, even though the format is different between them. A Notebook, which contains multiple pages analogous to spreadsheets (.xls) and workbooks (.xlw) in Excel, uses the extension .onepkg.

In Office 2010, the integration between Outlook and OneNote can really improve productivity. There's a Send to OneNote icon in the Outlook ribbon and an E-Mail Message icon in the Outlook section in the OneNote ribbon. Each one simplifies data transfer between the applications. I'll give you an example: I received an email from Lance Armstrong's LiveStrong foundation. If I select the Send to OneNote button with the LiveStrong foundation email selected in the Outlook Inbox (or use the same option when the email is opened in its own window), it will open the dialog box in which I can identify where the item should be saved. The default in Outlook is not to download images "to help protect your privacy." So the images were not downloaded. However, when you click the Send to OneNote icon, images accompanying an email are then retrieved from the server (or web server) and included in the OneNote page output.

With Outlook installed, OneNote can launch an email when you select Home, E-mail Page in the Office ribbon. This will place the OneNote page content into the body of a new email. Alternatively, if the recipient is also using OneNote, you can attach the OneNote page in OneNote format, with the .one extension. This is a global setting within OneNote. You can configure it by accessing File, Options, Advanced. Either way, Outlook automatically generates a new email message with the OneNote page name as the message subject and the default Outlook account as the sender.

When you send a OneNote page through the E-mail Page option in OneNote, you can include a disclaimer. By default, Microsoft adds the following text at the end of the email message: "Created with Microsoft OneNote 2010. One place for all your notes and information." In a standard HTML email, the boilerplate text is formatted in grey and is lighter than any black text within the message. This is configurable through the Advanced Options window. You can deselect the option to remove any additional text or customize the text to suit your company's needs. This option is good for adding disclaimers or support information to OneNote page content sent through Outlook.

If Outlook attaches the OneNote page as a .mht file, a MIME HTML web archive, then there is an odd issue where "the string value for Outlook 2010 is longer than any other program's string value in the RegisteredApplications key in the registry." This seemingly preventable problem is fixed by a simple registry entry or by installing a supplementary application. The options for remedying this issue are outlined in Microsoft knowledge base article 982991.

I've seen Microsoft MVPs use OneNote to store answers to frequently asked questions posed by people in forums and newsgroups. The MVPs typically save and configure standard responses to OneNote. Then they retrieve their canned reply by copying and pasting it to a web forum or to an email message. Help desks could use OneNote to store how-to information to send out as supporting documentation in response to

a help desk call. Human Resources might maintain basic forms that could be sent by email through OneNote as well.

—William Lefkovics

InstantDoc ID 125720

Q: How do I enable jumbo frames on an ESX v4.x server?

A: Jumbo frames are network frames that contain more than the standard 1500 bytes of payload, up to a maximum of 9000 bytes. (This size is sometimes called the Maximum Transmission Unit—MTU). For networks with a high degree of reliability, such as your internal LAN, jumbo frames' greater payload size can mean fewer network acknowledgements during data transfer, and thus better network performance for certain kinds of transfers, such as iSCSI storage traffic.

ESX v3.x included experimental support for jumbo frames on software iSCSI. In ESX v4.x, jumbo frames are fully supported. Enable jumbo frames for a vSwitch by increasing its MTU inside the ESX Service Console:

```
esxcfg-vswitch -m <MTU> <vSwitch>
```

You can verify the switch has been given the new MTU value with the command:

```
esxcfg-vswitch -l
```

You'll have to enable jumbo frames at each hop throughout your network, so it takes more work than just enabling them at your ESX server.

—Greg Shields

InstantDoc ID 125688

Q: What's the difference between Windows 7 Federated Search and the Enterprise Search feature available in Windows 7 Enterprise and Ultimate?

A: One of the new features in Windows 7 that's available in all SKUs is Federated Search, which lets a Windows 7 client search a remote data source that supports OpenSearch 1.1 then locally show the results of the search that are sent back using RSS or Atom. All the searching is performed on the remote server using its own index. The only requirement on the client is a small OSDX file (Open Search Descriptor XML file) that tells the Windows 7 client how to talk to the remote server and the URL to actually use for the remote communication.

You can install multiple OSDX files on a client. A great example of this in action is to add a Federated Search to your organization's SharePoint site. Once the federation is in place (through the OSDX file), a user can perform a search on their desktop and tell the search to use SharePoint as the search target. You can also add federated search for Internet services such as Amazon.com, making it easy to quickly find the latest DVDs.

Enterprise Search Scopes in Windows 7 Enterprise and Ultimate lets organizationally defined search targets (up to five) be targeted to the client desktops so the organization can help guide where users should focus their searches. These search locations appear automatically anywhere a search is performed, such as on the Start bar or in Explorer (as Search Again links). Organizations would likely push their Intranet locations and SharePoint sites as enterprise-wide search targets.

—John Savill

InstantDoc ID 125838

Q: Can I use PowerShell to create and compare a configuration baseline?

A: Almost anything you can "Get" is eligible for this trick. Start by getting your baseline into an XML file. Let's say you're working with services. You'd use the command

```
Get-WmiObject -class Win32_Service |
  Export-CliXML baseline.xml
```

Then, when you're ready to compare the current state to the baseline, you'd run:

```
Compare-Object (Import-CliXML baseline
  .xml) (Get-WmiObject -class Win32_
  Service)
```

That all gets written on one line. Compare-Object also has an alias, Diff.

The comparison will compare every attribute of those objects. That works well for many management objects such as services but not all. Processes, for example, are constantly changing. You expect a process's CPU usage, for example, to change, so you wouldn't compare everything. Instead, you'd just pick one or two properties to compare, as in

```
Compare-Object (Import-CliXML
  processes.xml) (Get-Process)
  -property name
```

—Don Jones

InstantDoc ID 125726

Q: What, exactly, is DNS name devolution? Are there any security risks linked to this DNS feature? Has anything changed in Windows 7 and Windows Server 2008 R2 to better protect my Windows platforms against these security risks?

A: DNS name devolution is a built-in feature of the Windows DNS Client. Users of Active Directory (AD)-joined computers can use DNS name devolution to connect to resources using an unqualified name, such as mailserver1, instead of using a Fully Qualified Domain Name (FQDN), such as mailserver1.emea.mydomain.net. The DNS name devolution feature allows the DNS client to automatically append portions of the AD-joined computer's primary DNS domain suffix (for example, "emea.mydomain.net") to the unqualified name during the DNS name resolution process.

For example, when a user on a computer that's a member of the emea.mydomain.net domain uses the resource name mailserver1, the DNS client will automatically try to resolve the mailserver1.emea.mydomain.net and mailserver1.mydomain.net FQDNs.

An important parameter in the DNS name devolution process is the devolution level. It refers to the number of labels in the primary DNS domain suffix at which the devolution process stops. Labels are the parts of a DNS name that are separated by dots. In the above example, emea, mydomain, and net are the three labels of the emea.mydomain.net domain suffix.

In Windows versions prior to Windows 7 and Windows Server 2008 R2, the DNS name devolution level is always two. This

means that if you type mailserver1 and the primary domain suffix is france.emea.dc.net, the DNS client will first try to resolve mailserver1.france.emea.dc.net, then mailserver1.emea.dc.net, then finally mailserver1.dc.net. At this point devolution will stop, because only two labels—dc and net—are left.

The default devolution level of two creates a security risk. It may cause domain-joined computers to connect to a malicious computer on the Internet that's outside of the control of an organization's AD domain. Let me illustrate this with an example.

A domain-joined computer's primary domain suffix is mycompany.fl.us (mycompany is located in Florida, hence the extension fl.us) and tries to connect to mailserver1. In this example, the DNS client will try to resolve mailserver1.mycompany.fl.us and mailserver1.fl.us. The last name in this list, mailserver1.fl.us, is outside of the control of my company. If a malicious person has registered mailserver1.fl.us in the DNS, the name resolution will succeed, the domain-joined computer will try to connect to it, and the malicious user could spoof an internal server.

In Windows 7 and Server 2008 R2, Microsoft changed the default DNS devolution behavior such that it cannot cause an internal client to connect to an external computer. Microsoft also provides an update for older Windows platforms to bring the new DNS devolution logic to these older platforms. Microsoft offers more information on this fix.

The DNS devolution logic has changed as follows: If the number of labels in the AD forest root domain's DNS name is one or a machine's primary DNS suffix doesn't end with the forest root domain's DNS name, DNS devolution is automatically disabled. For example, if a computer is a member of the mycompany.com domain and the forest root domain name is mycompany.fl.us, devolution is disabled (mycompany.com does not end with mycompany.fl.us).

If a machine's primary DNS suffix ends with the forest root domain's DNS name, the devolution level is automatically set to the number of labels in the forest root domain. For example, if the computer is a member of the research.mycompany.fl.us domain and the forest root domain name is mycompany.fl.us, the devolution level is set to three (which matches the number of labels in mycompany.fl.us).

You can enable name devolution from the DNS tab in the advanced properties of the TCP/IPv4 and TCP/IPv6 protocols of a Windows box's network interfaces. When you click Append primary and connection specific DNS suffixes and select Append parent suffixes of the primary DNS suffix, name devolution is enabled.

You can also centrally configure name devolution with the following Group Policy settings, which are located in the Computer Configuration\Administrative Templates\Network\DNS Client GPO container:

- Primary DNS Suffix Devolution: This Group Policy Object (GPO) setting controls the HKLM\SOFTWARE\Policies\Microsoft\Windows NT\DNSClient\UseDomainNameDevolution registry value.
- Primary DNS Suffix Devolution Level: This GPO setting controls the HKLM\SYSTEM\CurrentControlSet\Services\Dnscache\Parameters\DomainNameDevolutionLevel registry value.

—Jan De Clercq

InstantDoc ID 125767

Q: How and why would I increase the amount of RAM assigned to the ESX Service Console?

A: For most ESX 4.x environments, the default 358MB of RAM is enough to handle the needs of the Service Console. But in some situations, you may find that you need additional horsepower to get the job done. This can be because you've installed additional software to the ESX host or because you're running custom processes that require extra RAM.

To change the amount of RAM assigned to an ESX host's Service Console, first log into the vCenter Client and select the ESX host from the inventory. Click the Configuration tab, followed by Memory, and then Properties. In the resulting window, enter a new value for the amount of assigned RAM. This value must be between 256MB and 800MB. Click OK and reboot the host to complete the assignment.

—Greg Shields

InstantDoc ID 125718

Q: I'm trying to enable the Hyper-V role on a box that doesn't have the hardware required for Hyper-V. I just need the Hyper-V WMI parts for Virtual Hard Disk (VHD) mounting. What can I do?

A: If you try to enable the Hyper-V role through Server Manager or the PowerShell interface on a box that doesn't meet the hardware requirements, you'll receive an error that the hardware doesn't meet the requirements and the installation of Hyper-V will be blocked. A workaround is to use ocsetup, which will allow the role to be installed, but this isn't a Microsoft-supported action. Just run

```
ocsetup Microsoft-Hyper-V
```

You should set the hypervisor launch to disabled, because it won't be able to load anyway.

—John Savill

InstantDoc ID 125810

Q: How can I check to see if a registry key exists in PowerShell?

A: Keep in mind that the shell maps two drives—HKLM: and HKCU:—to the registry. Using those, you can use the Test-Path cmdlet to see if a key, which will behave more or less like a file system folder, exists.

—Don Jones

InstantDoc ID 125571

Q: Can I move a VM between Hyper-V hosts with different CPU types if the virtual machine (VM) is turned off?

A: Providing the VM is turned off (which also means it can't be in saved state), a VM can be moved between hosts with different versions of processors and even different processor manufacturers (AMD/Intel).

—John Savill

InstantDoc ID 125710

Kerberos in Active Directory

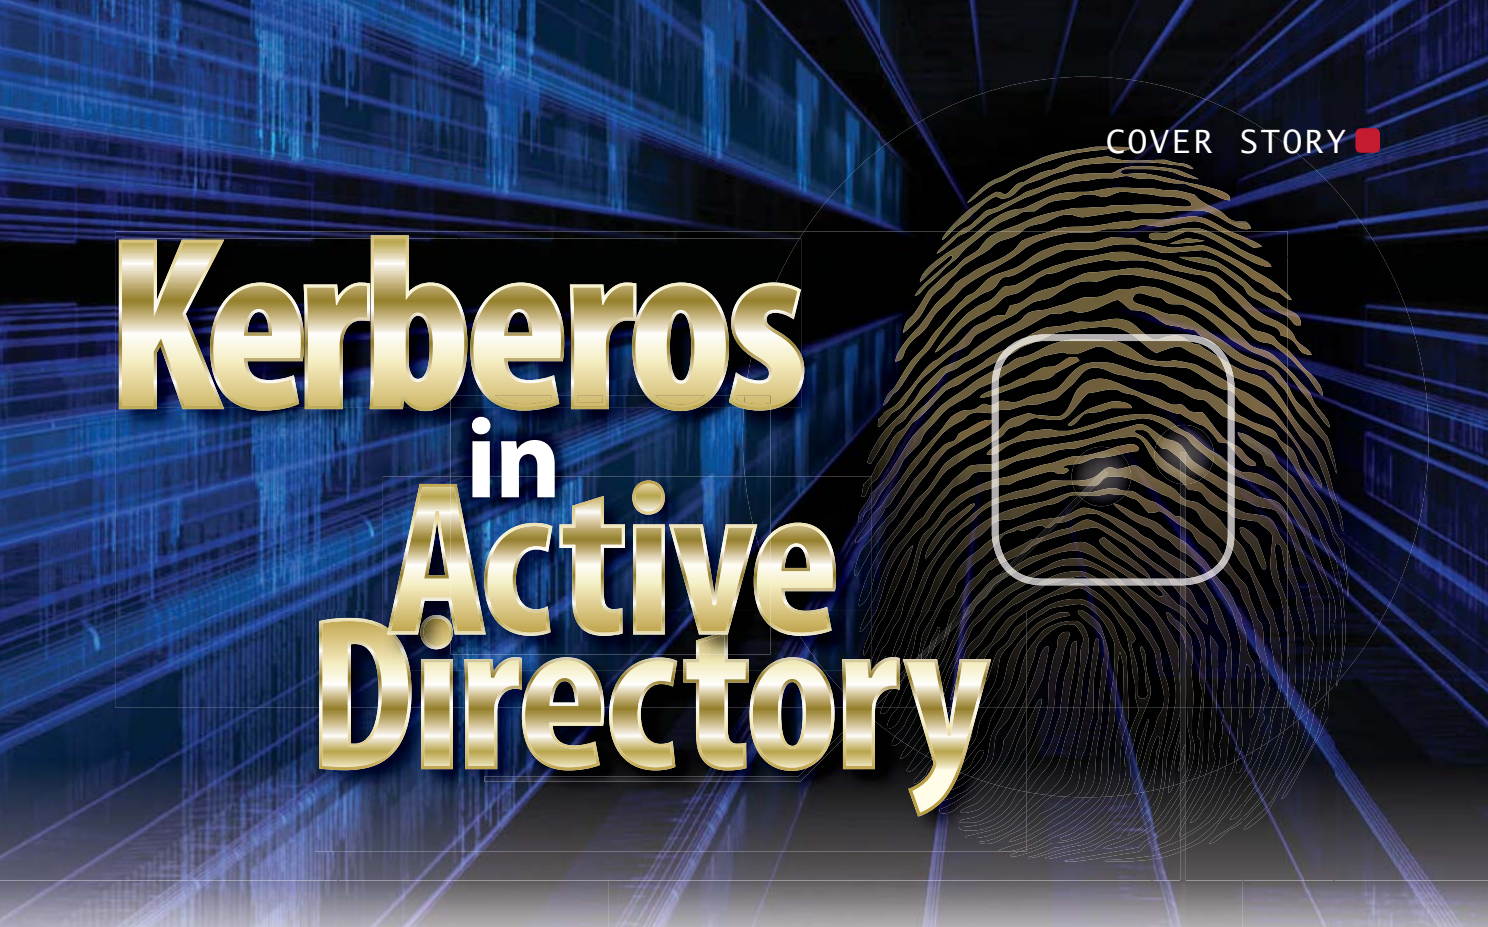

Revealing the underpinnings of AD authentication

by Brian Desmond

Kerberos might seem like black magic, but it's one of Active Directory's (AD's) key underpinnings. Most of Kerberos's configuration is abstracted, making interaction with the protocol uncommon. However, Kerberos is used every time you log on to an AD-joined machine, as well as when you access network resources such as file shares and applications.

Rather than transmit your actual password over the network, Kerberos operates with a series of tickets. At a high level, when you log on to a machine you generate a series of exchanges with the domain controller (DC) that, if successful, ultimately grants you a ticket-granting ticket (TGT). Each time you subsequently want to access a service, such as a file share or application, you use the TGT to apply for a service ticket to the service or application you want to access.

Authentication

When a machine authenticates during logon, it sends an AS_REQ or Authentication Service Request Kerberos message to the DC. In Kerberos, we call the DC a Key Distribution Center (KDC). Figure 1 shows the critical contents of such a request. The client name is either the user's user principal name (UPN) or the user's legacy username (sAMAccountName). The service name in this case is always the same, a request to the krbtgt service in the user's domain (or realm). To prevent replay attacks whereby an attacker recycles an AS_REQ message, the current time is encrypted using a hash of the user's password. This is where the five-minute clock skew tolerance comes into play. If the encrypted timestamp isn't within five minutes of the current time, the request is rejected.

When the KDC receives an AS_REQ message, the first thing it tries to do is decrypt the encrypted timestamp using its local copy of the user's password hash. If this operation fails, then an error is returned to the client and the request doesn't proceed any further. If the decryption is successful and the timestamp is within acceptable limits, the KDC returns an AS_REP (Authentication Service Reply) message to the user, with an embedded TGT. Figure 2 shows the contents of the AS_REP message. At this point, the user's machine caches the TGT and session key for the lifetime of the TGT and disposes of the user's password. By default, TGTs issued by AD KDCs expire after ten hours.

AD's Kerberos implementation requires the encrypted timestamp in the AS_REQ message. This initial exchange is known as pre-authentication. Kerberos as a standard doesn't require the encrypted timestamp and instead is perfectly happy with an AS_REQ message that simply contains the client name and service name. The problem with this approach is that an attacker could mount a dictionary attack because each request containing a unique client name would either return an error stating the client wasn't found or would yield a valid TGT for a given user.

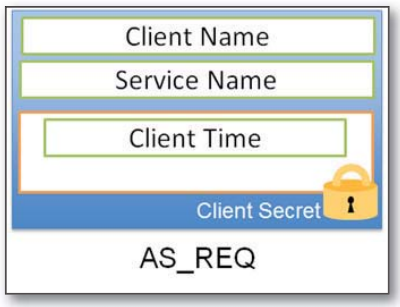

Figure 1: AS_REQ Authentication Service Request message

Figure 2: AS_REP Authentication Service Reply message

As Figure 2 shows, the AS_REP message contains two pieces of encrypted data. The first component is encrypted with a hash of the user's password and contains a session key and ticket expiry timestamp. The session key is used to encrypt future communication with the KDC. The second component, the TGT, is encrypted with the KDC's secret. The KDC's secret in AD's Kerberos implementation is stored as the password to the krbtgt user account that exists in every AD domain. The krbtgt account is created when the first DC in a domain is promoted; this account is crucial to the domain's operation.

Obtaining a Service Ticket

In Kerberos, anything you want to access is called a service. Services include file and print servers, database servers, and internal web applications. To access a service, you present a service ticket to the service. The first step is to identify the service principal name (SPN) of the service you want to access. Your machine or the application involved is responsible for forming the SPN.

In the case of a file server called srv01.contoso.com, the SPN for the file service would be cifs/srv01.contoso.com. Likewise for a web application at web01.contoso.com, the SPN might be http/web01.contoso.com. A common problem that AD administrators encounter is one of duplicate SPNs, in which more than one user or computer in the directory has the same SPN registered in its servicePrincipalName attribute. When this happens, the KDC doesn't know how to respond to a service ticket request because multiple services with the same name exist.

If you look at the servicePrincipalName attribute in AD, you'll notice that neither of the two sample SPNs mentioned are defined on any computer account in AD. To reduce the amount of duplicate data stored in the directory, AD maintains a forest-level mapping of default SPNs to every computer account in the directory. Literally dozens of built-in SPNs apply to every computer defined in the attribute. The data is stored in the spnMappings attribute of the Directory Service object in the Configuration partition (under CN=Directory Service, CN=Windows NT, CN=Services, CN=Configuration); this attribute maps each computer object's HOST service to a variety of other services. These services include the CIFS and HTTP services. If you wanted to define a particular service on every machine in your forest, you could easily modify this attribute to add the service to the list of predefined services.

Figure 3 shows the contents of a TGS_REQ (Ticket Granting Service Request) message that the KDC receives from a client when the client tries to access a service. The first piece of information is the SPN of the service the client is requesting a ticket for. Encrypted with the session key the client received as part of the earlier AS_REP message is the client's name and a timestamp. This information is again used to prevent replay attacks whereby an attacker reuses a request message. Also included is a copy of the TGT the client received earlier.

When the KDC receives a TGS_REQ message, if one entry for the SPN is specified, the timestamp is within range, and the TGT is valid (and unexpired), the client receives a service ticket as part of a TGS_REP message, as Figure 4 shows. The TGS_REP message contains everything the client needs to access the service.

Encrypted with the session key in the AS_REP message that Figure 2 shows is a second session key for communicating with the service the client has a ticket for. Encrypted with the service's secret (e.g., the password of the machine account or service account) is a service ticket the client will cache and use whenever it needs to access the service. Much like the TGT, service tickets have a maximum lifetime for reuse (ten hours by default in AD's implementation of Kerberos). With a service ticket in hand, the client can now request access to the service.

Accessing Services

After the client has a service ticket, the application accessing the service can present that ticket to the service and request access. The mechanics of presenting the service ticket aren't nearly as standardized as for obtaining the ticket because every application is different. In the case of an HTTP service, the service ticket is embedded in the headers of the HTTP request.

The service ticket is presented to the application in the form of a Kerberos AP_REQ message, as Figure 5 shows. The service decrypts the service ticket and obtains the session key, which it can use to decrypt the timestamp and client name fields, which are in turn used to validate the authenticity of the service ticket. Even if the service accepts the service ticket, at this point the client has merely authenticated to the service. The task of authorization is still up to the service, based on the information it has about the client.

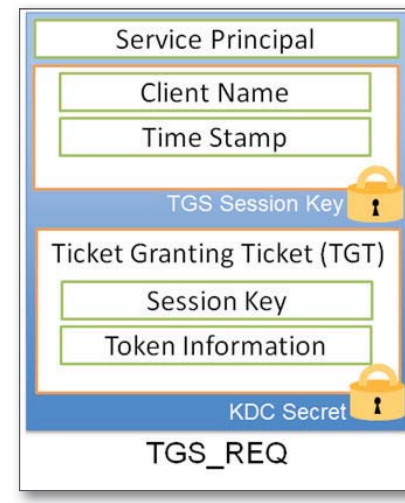

Figure 3: TGS_REQ Ticket Granting Service Request message

Smarter technology for a Smarter Planet:

It's time to ask smarter questions.

What exactly does a benchmark mean? For the last five years, IBM DB2® on Power Systems™ has ranked first on three of the industry's leading performance benchmarks, longer than Oracle and Microsoft combined.[1] But shouldn't we be asking our IT to deliver more than just raw performance? What really matters isn't some abstract measure of performance, it's what companies actually do with it. For instance, Coca-Cola Bottling Company is using DB2 on Power to reduce licensing, maintenance and storage fees by $350,000. EuResist is using an integrated analytics solution to predict the most effective drug combinations for individuals with HIV, with 78% accuracy. And the Dubai Gold & Commodities Exchange is working with IBM Security Services to achieve system uptime of over 99.9%. On a smarter planet, these are the benchmarks that matter.

A smarter business is built on smarter software, systems and services.
Let's build a smarter planet. ibm.com/questions

1- Based on number of days of performance leadership for the IPC-C, TPC-H 10TB, and SAP 3-Tier SD benchmarks between June 1, 2005, and June 1, 2010. For more information, see http://www.tpc.org and http://www.sap.com/solutions/benchmark. TPC, TPC-C and TPC-H are trademarks of the TPC. IBM, the IBM logo, ibm.com, DB2, Power Systems, Smarter Planet and the planet icon are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml. © International Business Machines Corporation 2010.

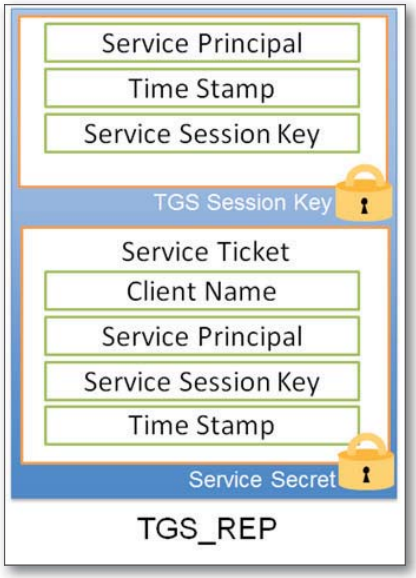

Figure 4: TGS_REP Ticket Granting Service Reply message

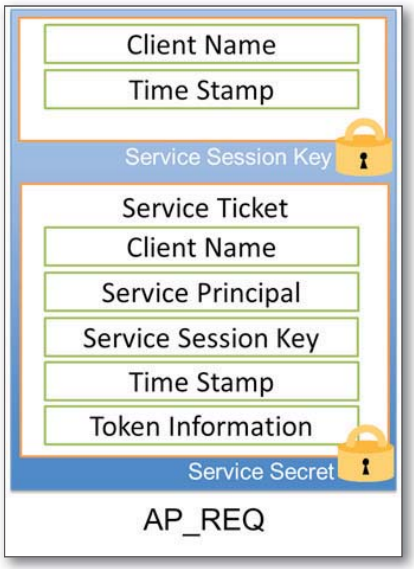

Figure 5: AP_REQ Application Request message

The service ticket typically also includes data known as the Privilege Attribute Certificate (PAC). In Figure 5, the PAC is called *Token Information.* This is the same token information the KDC included in the user's TGT (see Figure 2). The PAC is composed of information such as the user's SID, group membership information, and user security rights/privileges. When a user presents a TGT to the KDC to request a service ticket, the KDC copies the token information from the TGT and includes it in the service ticket's PAC field. This is the information the service uses to construct an access token for the user and to verify the user's authorization, typically based on group membership.

An additional Kerberos message known as an AP_REP or Application Reply is permissible after the user presents a service ticket in the AP_REQ message that Figure 5 shows. The Application Reply message is optional; in general, the application won't send such a message unless an error occurs. One example of when an AP_REP message would be generated is in the case of a client that requests (in the AP_REQ message) that a service prove its identity through a process known as mutual authentication.

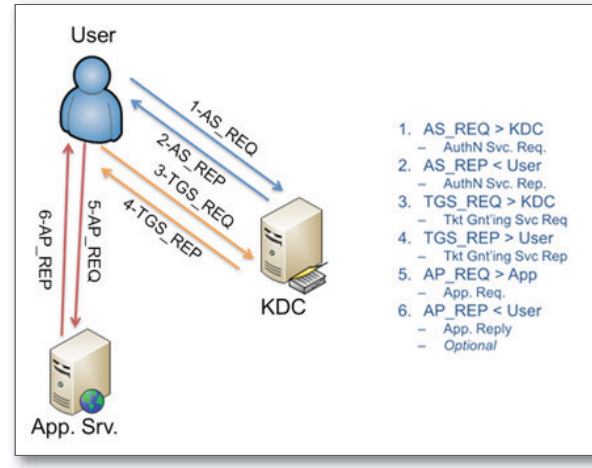

Figure 6: Kerberos authentication and service access message flow

Process Overview

Figure 6 shows a recap of the message flow when a user decides to access a service on an application server. The user logs on to his or her workstation, generating an AS_REQ and AS_REP message sequence to the KDC, where the user receives a TGT if the credentials are valid. The user's TGT is subsequently cached in memory and each time the user wants to access a service (e.g., file server, print server, web application), the user presents the TGT back to the KDC and requests a service ticket for the service. The user receives the service ticket and presents it to the application to request access.

If you're using read-only domain controllers (RODCs), the message flow is a bit different from what Figure 6 shows, depending on your password replication policies and where the passwords are cached. Remember that the KDC needs access to the service's and user/computer's secrets to issue service tickets and TGTs, respectively. By default, an RODC caches no passwords, so it has no access to the necessary secrets.

If an RODC receives a request for a TGT or service ticket that it can't fulfill (because of an uncached password), it passes the request back to a writeable DC that holds the necessary secrets to encrypt the response. The writeable DC sends the response to the RODC, which in turn sends it back to the client. If the RODC has the password cached for either the user or the service the user is trying to access, the RODC simply issues the TGT and/or service ticket in the same manner as shown in Figure 6.

If you refer back to Figure 2, you'll see that the user's TGT is encrypted using the KDC secret. This secret is the password to the krbtgt account, which all AD domains have. Considering the baseline design assumption of an RODC (i.e., the fact that it's compromised by default), it would be potentially catastrophic if the RODC contained the password to the domain-wide krbtgt account. If the RODC were compromised, the attacker could issue TGTs independently. To mitigate this risk, each RODC has its own krbtgt account and is never aware of the passwords for either the domain-wide krbtgt account or another RODC's krbtgt account. Writeable DCs have access to the passwords for all the RODC krbtgt accounts so they can decrypt TGTs issued by any RODC in the domain.

Transparent Authentication

Kerberos is almost transparent to AD administrators. Unlike many authentication schemes, Kerberos never requires that a password traverse the network, nor does it require a user's password to be stored in memory after the initial logon. Both these benefits greatly reduce the inherent risk of other authentication protocols.

InstantDoc ID 125786

Brian Desmond
(brian@briandesmond.com) is a Directory Services MVP and senior consultant for Moran Technology Consulting in Chicago. Brian is author of *Active Directory,* 4th edition (O'Reilly), and blogs at www.briandesmond.com.

HOB
Secure Business Connectivity

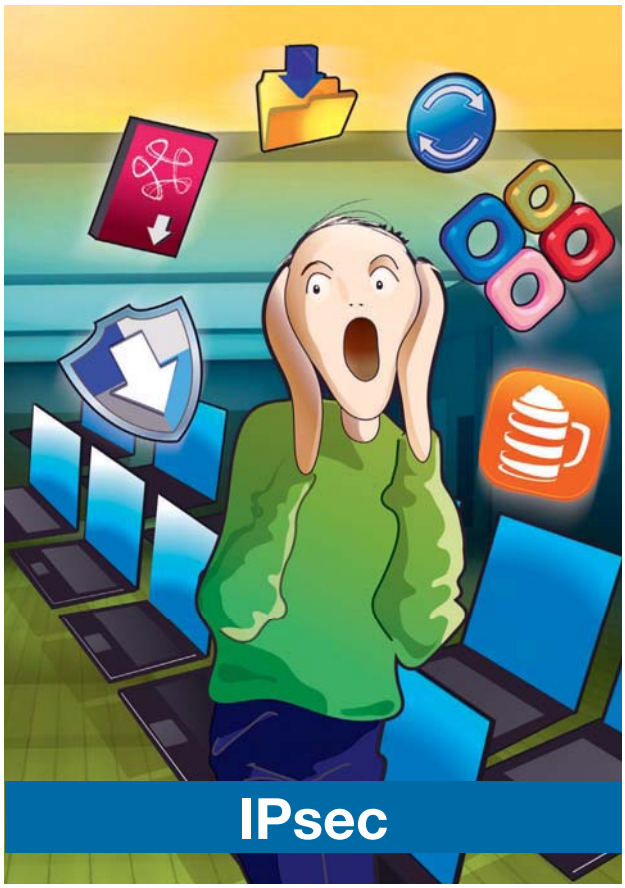

HOB RD VPN Gives You Peace of Mind

HOB RD VPN

Secure Remote Access and No Installation

Unlike conventional SSL VPNs, as well as IPsec VPNs, HOB RD VPN provides a browser-based secure remote access solution. With HOB RD VPN, enterprise administrators no longer need to be occupied with the time-consuming installation of drivers or software on client syems. Users are presented with a mobile, platform-independent virtual desktop pre-loaded with all required applications – with all data securely stored in the data center.

HOB PPP Tunnel (Fully Replaces IPsec)

HOB PPP Tunnel's driverless technology provides users with full network access from any Microsoft Windows Vista, Windows 7, Apple Mac OS X or Linux system. With two-factor authentication, SSL encryption up to 256-bit AES, and client integrity check – PPP Tunnel delivers all the security of an IPsec solution, without compromising performance and ease-of-use.

HOB RD VPN has all the required connectivity clients (e.g., HOBLink JWT), which are run on the client machine without having to be installed there.

The **HOB RD VPN** Security Suite has been certified by the **BSI (German Federal Office for Information Security)** in acc. w. the Common Criteria.*
This shows that with **HOB**, remote access truly is secure!

FEATURE

Advanced Active Directory Security

Enumerate accounts protected by AdminSDHolder and fix resulting ACL problems

by Russell Smith

Back in the days of Windows 2000, I was asked to design a delegation model for an Active Directory (AD) project. The project required that a small set of trusted users be granted the right to reset passwords on accounts in an organizational unit (OU), which also belonged to the Domain Admins group. At the time this didn't appear to be a particularly difficult task. However, I was surprised to find that ACLs that had been applied to the accounts for the purposes of delegating appropriate permissions had disappeared. My initial reaction was that there must be something wrong with AD in the lab or that someone had changed the permissions back to their defaults. So we reset the ACLs in the hope that they'd stick, only to find that they had disappeared again.

A little more investigation into the problem revealed that this behavior is by design. A security mechanism called AdminSDHolder defines the ACLs applied to a list of protected users and groups and by default prevents them from inheriting permissions from their parent object in AD.

A process called SDProp (Security Descriptor Propagator) runs hourly to enforce ACLs on objects protected by AdminSDHolder. So, using the wizard in the Active Directory Users and Computers (ADUC) management console to delegate the reset password rights for accounts in an OU might not be effective if any of the child objects are members of a group protected by AdminSDHolder.

Users and Groups Protected by AdminSDHolder

In Win2K, AdminSDHolder protected just four key groups: Administrators, Domain Admins, Enterprise Admins, and Schema Admins. Over the years, that number has been expanded. Now in Windows Server 2008 and Server 2008 R2, AdminSDHolder protects the following users and groups: Account Operators, Administrator, Administrators, Backup Operators, Domain Admins, Domain Controllers, Enterprise Admins, Krbgt, Print Operators, Read-only Domain Controllers, Replicator, Schema Admins, and Server Operators.

Accounts that belong to a protected group should be treated as special-purpose accounts. For example, a user account that's a member of the Backup Operators group should be used only for performing backup and restore operations.

When a user becomes a member of a protected group, the user no longer inherits permissions from its parent object in AD. Additionally, some default ACLs are removed and others added.

Accounts added to protected groups have different ACLs than standard user accounts in AD. This can make an impact on functionality, such as the inability to upload a user's certificate to the Global Address List (GAL) in Exchange due to missing SELF permissions on the user's AD account.

Working with the adminCount Attribute

AdminSDHolder determines whether a user object should be protected by enumerating the user's group membership, including nested groups. If a user is deemed to belong to a protected group,

it's stamped with the ACLs as set on the AdminSDHolder object in AD. The user's adminCount attribute is also set to 1. You can run an LDAP query against AD to determine which user accounts or groups are protected.

The only surefire way to determine if an account is protected by AdminSDHolder is to expand its transitive group membership. That's because the adminCount attribute isn't set back to zero when an account is removed from a protected group.

In Server 2008 R2, you can run an LDAP query using Windows PowerShell's new AD administration cmdlets. Start by importing the AD module into PowerShell so that you can use the AD cmdlets. The Get-ADUser command below will list all user objects where the adminCount attribute is set to 1:

```
import-module activedirectory
get-aduser -ldapfilter `
  "(objectcategory=person) `
  (admincount=1)"
```

This Get-ADGroup command does the same, but for AD groups:

```
get-adgroup -ldapfilter `
  "(objectcategory=group) `
  (admincount=1)"
```

If you're not using Server 2008 R2, you can use the LDAP Data Interchange Format Data Exchange (LDIFDE) tool, where -f specifies a filename for recording the query results and -r specifies the LDAP filter:

```
ldifde -f ldifdeoutput.txt -r
  "(&(objectcategory=
  person)(objectclass=user)
    (admincount=1))"
```

You should also bear in mind that the adminCount attribute might be set to 1 even if a user isn't a direct member of a group protected by AdminSDHolder. This is because AD also sets the adminCount attribute to 1 if a user has indirect membership of a protected group through nested groups.

Clearing the adminCount Attribute

When a user is removed from a protected group, the adminCount attribute on that person's user object stays set to 1. The account remains in a strange state of limbo in that it no longer receives ACLs from the AdminSDHolder object.

The account also doesn't inherit permissions from its parent container. Any ACLs previously applied from AdminSDHolder remain in place.

Setting the adminCount attribute to 1 isn't enough to protect a group or account with AdminSDHolder, so resetting the value to zero might not be strictly necessary. However, to be consistent, you should set the attribute back to zero.

Accounts that are demoted from a protected group should be deleted as a security precaution. This should ensure that any back doors created using the account can't be exploited once it has been demoted.

If you must keep the account, you should review the ACLs on the object and manually change the adminCount attribute to zero using ADSI Edit. To do so, follow these steps:

1. Log on to a Server 2008 R2 domain controller (DC) as a Domain Administrator. Find the ADUC console by searching for the phrase active directory in the Start menu's *Search programs and files* box.
2. In the ADUC console, click the View menu and make sure that Advanced Features is selected.
3. In the left pane, expand your AD hierarchy to find the container or OU with the user object you want to reset.
4. Right-click the object in the right pane and select Properties from the menu.
5. On the Attribute Editor tab, you'll see the adminCount attribute listed (see Figure 1), which you can edit if required.

When a user account is removed from a protected group, resetting the adminCount attribute isn't enough to enable the object's inheritance flag, which is the default setting for users not protected by AdminSDHolder. You'll need to change it manually as follows:

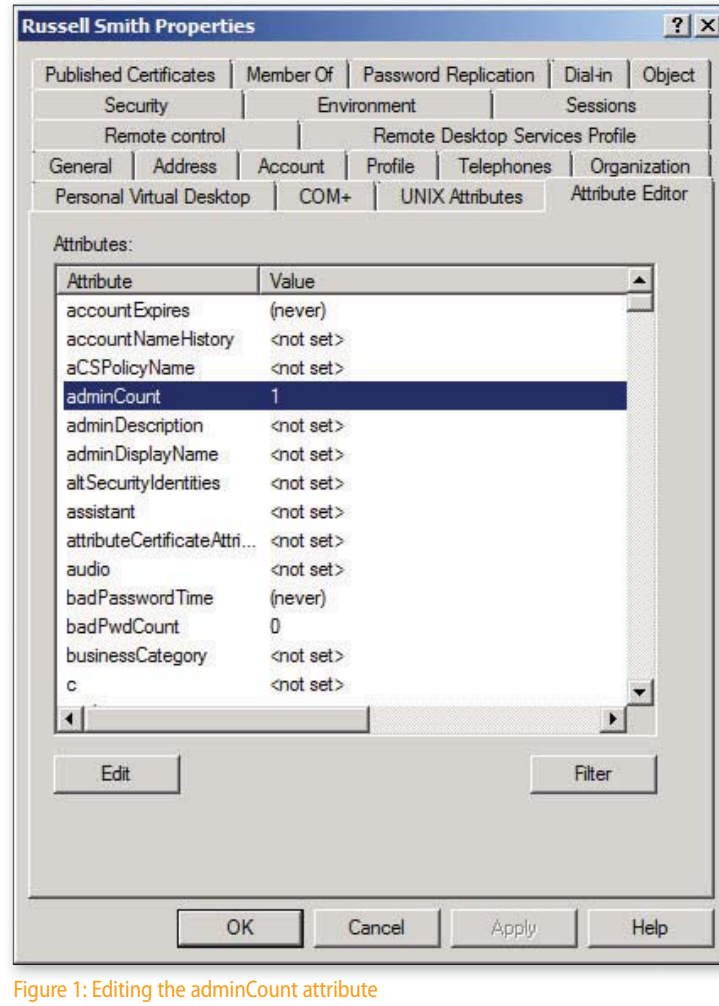

Figure 1: Editing the adminCount attribute

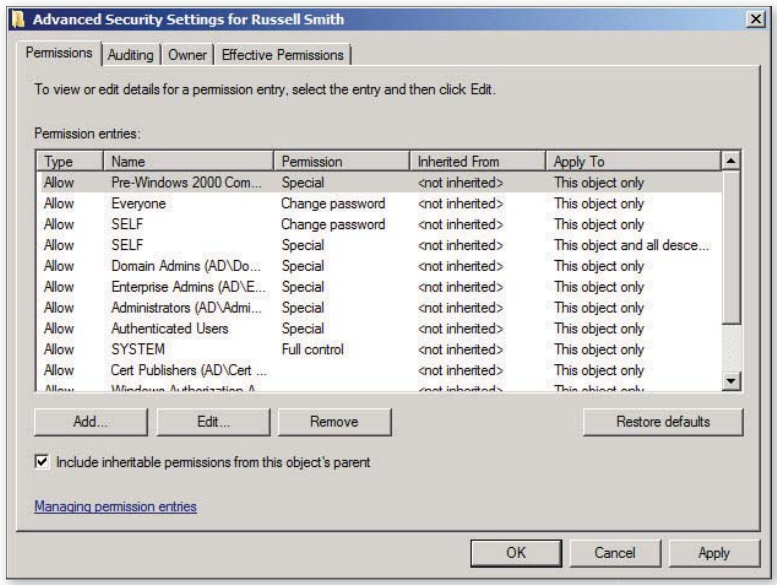

Figure 2: Enabling inheritance

1. Continuing from where we finished in the previous set of instructions, switch to the Security tab and click Advanced.
2. Select *Include inheritable permissions from this object's parent* and click OK (see Figure 2).
3. Click OK in the user's Properties dialog box to complete the procedure.

You can also reset the *Include inheritable permissions* flag and restore permissions that were assigned to the object when it was first created, as defined in the AD schema. To do this, click *Restore defaults* in the Advanced Security Settings dialog box. Be aware that *Restore defaults* is not the same as restoring ACLs from a known point in time.

You can modify security on the AdminSDHolder object in AD, so that the inheritance flag is enabled. Or you can modify security by applying different ACLs to all protected objects. Changes made to ACLs or the inheritance flag on the AdminSDHolder object itself are applied to the objects it protects.

To change permissions on the AdminSDHolder object in a lab, follow these steps:

1. Log on to a Server 2008 R2 DC as a Domain Administrator, and find the ADSI Edit console by searching for it in the Start menu's *Search programs and files* box.
2. In the ADSI Edit console, right-click ADSI Edit in the left pane and select *Connect to* from the menu.
3. In the Connection Settings dialog box, leave all the default settings in place and click OK.
4. In the left pane, double-click *Default naming context* to expand the directory, and drill down to CN=System.
5. Select CN=System in the left pane, right click CN=AdminSDHolder in the right pane and select Properties from the menu.
6. In the CN=AdminSDHolder Properties dialog box, switch to the Security tab and click Advanced.
7. In the Advanced Security Settings for AdminSDHolder dialog box, check *Include inheritable permissions from this object's parent* and click Apply.
8. Click Add, then browse for an AD group to which you want to permanently assign change and reset password permissions on protected objects. In this example, I am using a group called On Duty Security Team.
9. After you've successfully selected a group, click OK.
10. In the Permission Entry for AdminSDHolder dialog box (see Figure 3), select Descendant User objects from the *Apply to* menu.
11. In the Permissions list, check *Change password* and *Reset password*, then click OK.
12. In the Advanced Security Settings for the AdminSDHolder dialog box, click OK.
13. To complete the procedure, click OK in the CN=AdminSDHolder Properties dialog box.

The next time the SDProp code runs, any user objects protected by AdminSDHolder will have additional ACL entries. My example group, the On Duty Security Team group, can change and reset passwords. The inheritance flag will also be enabled so ACLs are propagated from the parent OU.

Forcing SDProp to Run

The SDProp code runs hourly, by default, on DCs that hold the PDC Emulator role. It enforces ACLs on groups protected by AdminSDHolder.

If you need to apply these settings immediately, the SDProp code can be forced to run. Use the following LDAP operation:

1. Log on to a Server 2008 R2 DC as a Domain Administrator, and find the LDP console by searching for it in the Start menu's *Search programs and files* box.

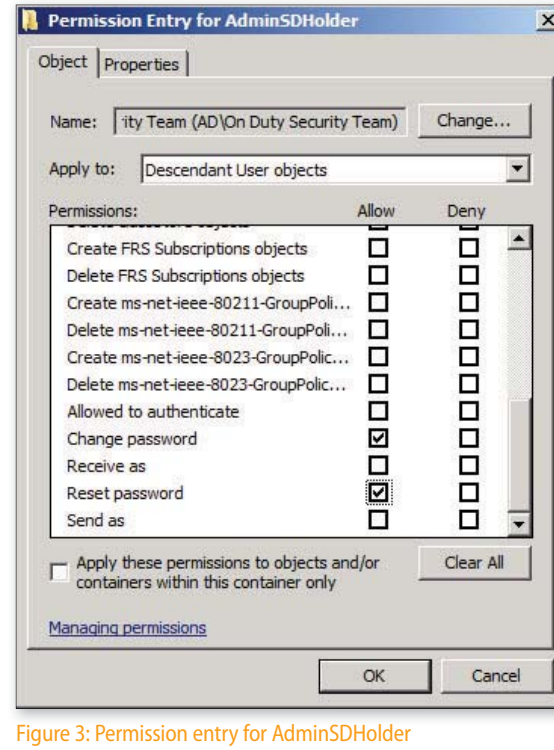

Figure 3: Permission entry for AdminSDHolder

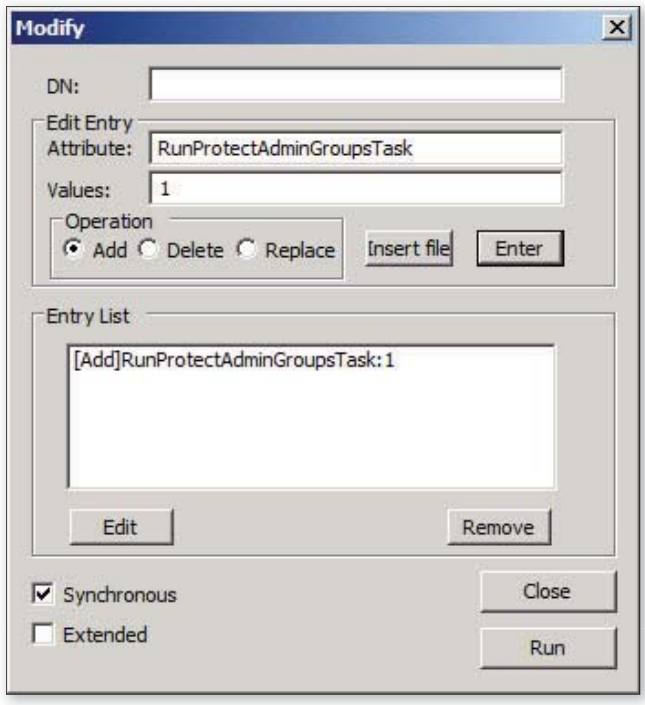

Figure 4: Modifying the RunProtectAdminGroupsTask

2. In the LDP console, select Bind from the Connection menu.

3. In the Bind dialog box, leave the default settings, and click OK.

4. Select Modify from the Browse menu. In the Edit Entry Attribute field, type RunProtectAdminGroupsTask, then type 1 in the Values field and click Enter (see Figure 4). For Windows Server versions prior to Server 2008 R2, replace the RunProtectAdminGroupsTask with FixUpInheritance and 1 with Yes.

5. Click Run, then Close. The update process may take some time depending on the size of your AD database.

Excluding Account, Server, Print, or Backup Operators from AdminSDHolder

As of Windows Server 2003 SP2, it's possible to exclude Account, Server, Print, and Backup Operators from AdminSDHolder protection. These operator groups are protected because they have privileges.

These privileges, such as log-on locally and shutdown, can have a serious effect on server security if used by the wrong person. If you must remove any of these groups from AdminSDHolder protection, you should consider removing their user rights. To do so, follow these steps:

1. Log on to a Server 2008 R2 DC as a Domain Administrator and open the ADSI Edit console.

2. In the left pane of ADSI, right-click ADSI Edit and select *Connect to* from the menu.

3. Select Configuration from the *Select a well known Naming Context menu* and click OK.

4. In the left pane, expand Configuration, CN=Services, CN=Windows NT. In the right pane, right-click CN=Directory Service, and select Properties from the menu.

5. In the CN=Directory Service Properties dialog box, select dsHeuristics on the Attribute Editor tab and click Edit.

6. In the String Attribute Editor dialog box, type 000000000100000f to exclude all four operator groups and click OK.

7. Click OK in the CN=Directory Service Properties dialog box and close ADSI Edit.

Each operator group has a binary value. The binary value for Account Operators is 0001 and Server Operators is 0010. For Print Operators, it's 0100, and for Backup Operators, it's 1000. The binary value must be converted to hexadecimal and added to the last position in the string ("f" in step 6 above).

If you want to exclude more than one group, add the binary values of the groups together. Then convert the result to hex.

For example, to exclude Print Operators and Backup Operators only, the dsHeuristics string would be 000000000100000c. You can use Windows calculator in Programmer mode to add binary numbers and convert them to hex.

Is this Really Best Practice?

As we've seen in this article, it's possible to modify the default ACLs for objects protected by AdminSDHolder. However, I wouldn't recommend changing AD default settings unless you have a very good reason to do so.

Modifications made to ACLs by AdminSDHolder can affect standard functionality for some key applications. These applications include Exchange, AD delegation, and any other AD function or application that relies on permission inheritance, such as BlackBerry servers.

Changing default AdminSDHolder behavior should be restricted. For example, you might restrict it to situations where political reasons dictate a specific solution.

If you follow best practices, changing AdminSDHolder functionality shouldn't be necessary. If you find you need to delegate permissions to objects protected by AdminSDHolder, it might be time to rethink your AD design to separate privileged tasks to dedicated user accounts.

Firecall accounts—special-purpose accounts used for server maintenance when elevated privileges are required—should be created for support purposes, added to a protected group, or assigned the required privileges. They should be disabled when not in use.

If accounts must be reused after they've been removed from a protected group, be sure to implement a procedure so that inheritance is re-enabled, ACLs are checked, and the adminCount attribute is set back to zero.

The most common reason to add a user account to one of AD's protected groups is so that a user can perform maintenance tasks on a DC. For more information on server access for support staff, see "Managing Privileged Access to Servers," June 2010, InstantDoc ID 104709.

If you can't escape the need for permanent access to servers hosting AD, read-only domain controllers (RODCs) allow administrative rights to be separated from access to AD.

InstantDoc ID 125777

Russell Smith
(rms45@rsitc.com) is an independent IT consultant specializing in systems management and security, and author of *Least Privilege Security for Windows 7, Vista and XP* (PACKT).

PROBLEM:

You need to configure a multiply redundant environment that can survive the loss of a network connection.

SOLUTION:

Add a new VMware ESX server to your cluster, and create a new iSCSI LUN that this server will use for VM storage.

WHAT YOU NEED:

An ESX server with two storage NICs; two Windows servers with Starwind Software's iSCSI SAN software installed

SOLUTION STEPS:

1. Configure and expose the StarWind SAN LUN on both Windows servers.
2. Configure the VMware ESX server's storage NICs.
3. Create a 1:1 mapping between the VMware ESX storage NICs and VMkernel ports.
4. Connect the VMkernel ports to the iSCSI initiator.
5. Connect the NICs to the StarWind SAN LUN and create a datastore.

DIFFICULTY:

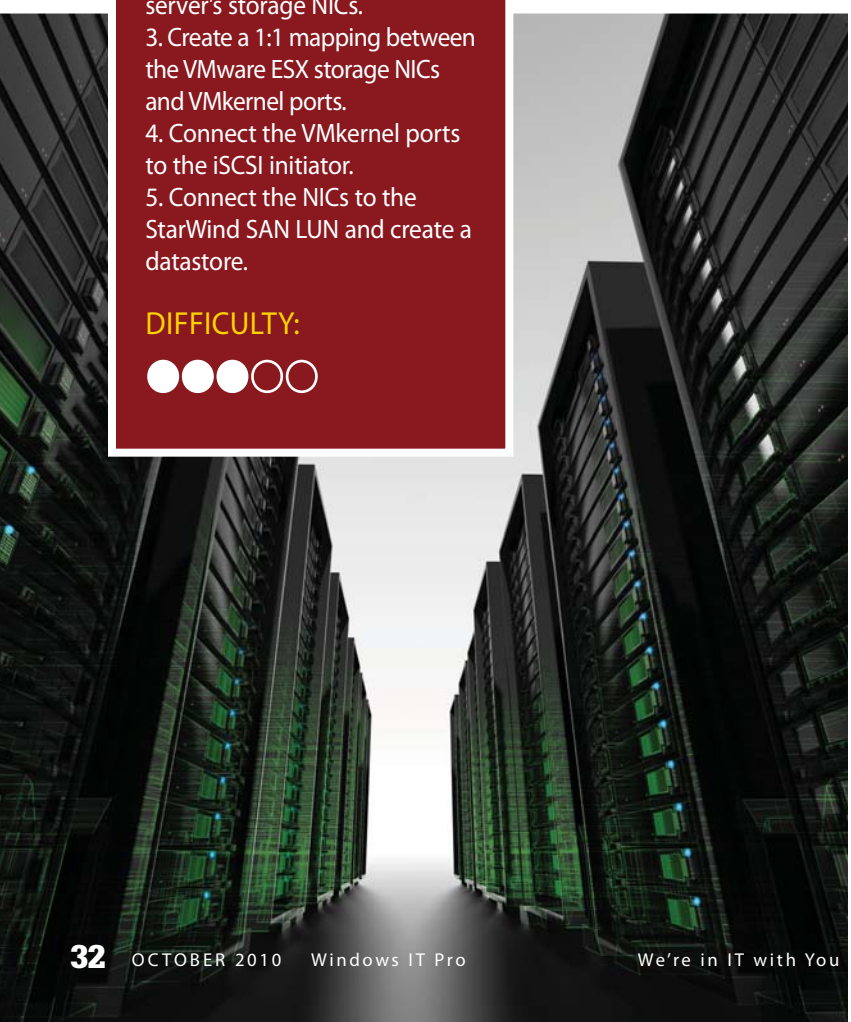

Configuring Highly Available iSCSI Storage for VMware ESX Server 4.x

The process should be easier than it is! Here are the step-by-step instructions you need

by Greg Shields

You have experience with VMware ESX Server. Who doesn't? But suppose you're tasked with the exciting job of adding a new VMware ESX server to your cluster. On top of that, you need to create a new iSCSI LUN that this server will use for virtual machine (VM) storage. Although you use VMware ESX Server every day to administer your VMs, you don't build new VMware ESX servers very often, and you're rusty on the skills necessary to connect a new server to your iSCSI SAN. The process isn't necessarily challenging, but some of the steps aren't obvious—and completing them in the correct order is important. This article will help get you going.

The Sample Environment

Before getting into the step-by-step instructions, let's take a look at a sample environment. Figure 1 shows a graphical representation of an environment in which two servers are running StarWind Software's iSCSI SAN software. I'm using StarWind's solution in my example in this article, but understand that every iSCSI SAN will offer a unique management console. The experience will be different, but the steps will be similar. (You can download a copy of StarWind's product at www.starwindsoftware.com. It installs on any Windows server and will give you a general idea of how SAN configuration works.)

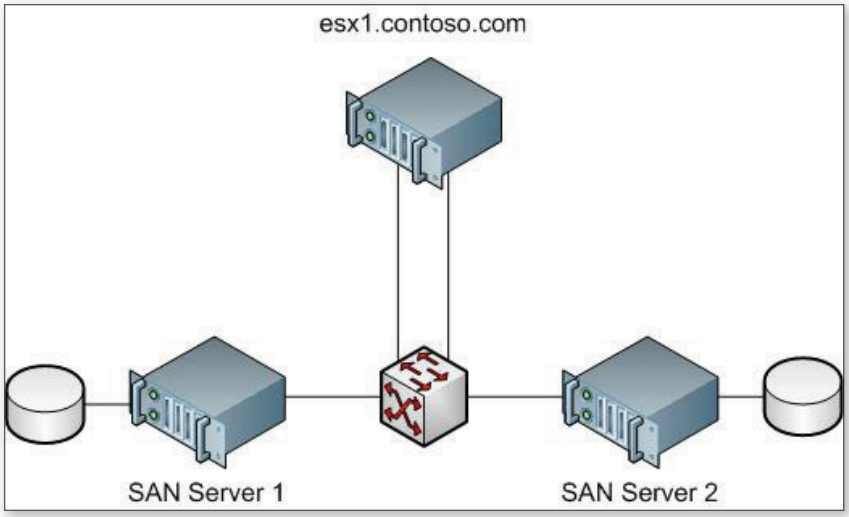

Figure 1: Our sample VMware ESX environment

Configuring the LUN

You need to create a redundant iSCSI LUN, so you'll need to create a LUN that's mirrored between the two SAN servers.

1. In the StarWind Management Console, ensure that you've added and connected to both hosts. If you're using the trial version of StarWind, the default logon and password are *root* and *starwind*, respectively.

2. Right-click Targets, and select Add Target. In the resulting screen, you'll be asked to provide a Target Alias and Target Name. The Target Alias is the friendly name for the iSCSI LUN you intend to create and is generally used only on the SAN device. The Target Name will be the iSCSI Qualified Name (IQN) used for the server-to-storage connection; it's the name you'll be seeing inside VMware ESX. You can safely leave the check box next to Target Name blank, allowing StarWind to create that IQN for you.

3. Set *Storage type* to Hard Disk, then set *Device type* to Advanced Virtual. Also, select *High Availability device*. These high-availability options will be available if you're using the Enterprise HA edition of the StarWind SAN software.

4. Creating a highly available LUN between two servers requires configuring that LUN with two partners. The first partner will be the server you used for the initial target creation. The second server is configured as a partner server. In this screen, provide the host's name as well as a username and password. You can use the same default username and password you used earlier, as long as they haven't been changed.

5. Creating this partner connection requires creating a Partner Target Alias and Partner Target Name, which must be different than the name in the earlier step. By default, the console will append the word Partner to your previous name.

6. You need to define the location where StarWind will store the file that will eventually become your VMware ESX LUN. You'll also need to configure how large you want that LUN to be.

7. You also need to configure the data synchronization channel parameters, which are effectively the target's network settings for the LUN. Figure 2 shows that an interface is configured for both partners by IP address. You can identify which partner is primary versus secondary, as well which port number is used for iSCSI traffic.

8. The StarWind Management Console next asks how you want to initialize the disks. Because these are brand-new disks, select *Clear virtual disks*. Click Next through the following series of screens, then Finish to complete the LUN creation. StarWind will require a number of minutes to synchronize between the two servers. Allow this process to complete. You'll know the process is over when the yellow or green warning symbols next to each target have disappeared.

Remember that your SAN software will have a slightly different series of steps for configuring this LUN. However, these sample steps are useful for setting up a demonstration environment if you don't yet have an iSCSI SAN or if you're still learning.

SAN uptime is absolutely important in VMware ESX environments. All of VMware's high-availability features work great—but only if your SAN never goes down. StarWind Software's Enterprise HA

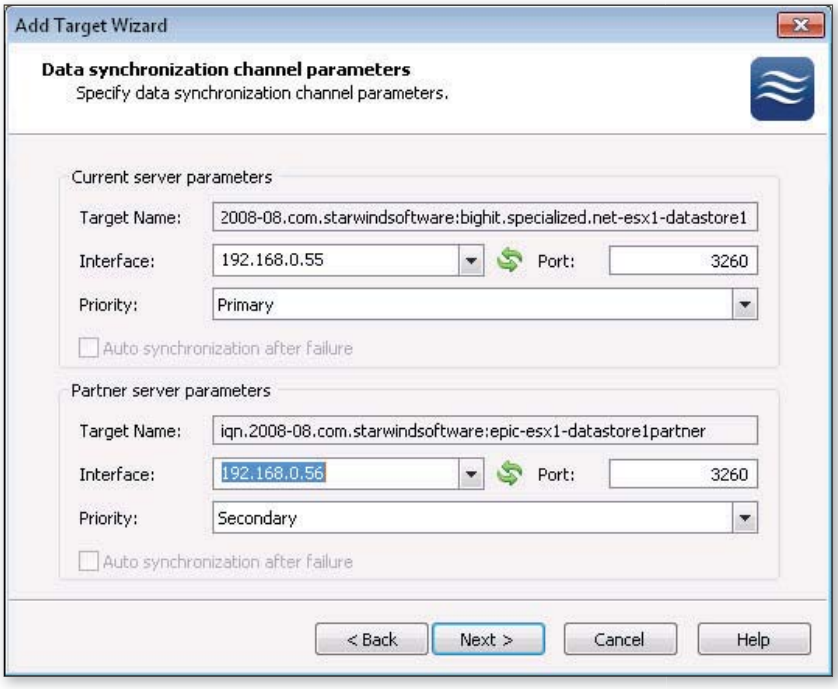

Figure 2: Configuring data synchronization channel parameters

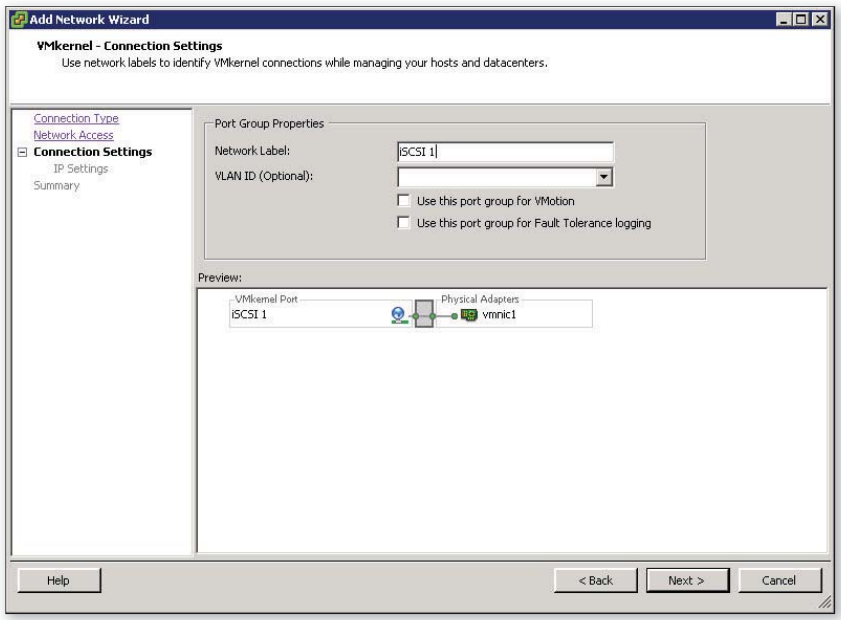

Figure 3: Configuring VMkernel connection settings

edition maintains that always-on SAN through data replication. Using it, your virtual environment can survive the loss of either SAN server without causing running VMs to go down.

Configuring Storage NICs

At this point, you've completed half of the connection's configuration. The first half created the LUN and prepared the iSCSI target for a connection from the VMware ESX server. The second half involves configuring the connection on the VMware ESX server itself.

So, your next task is to configure two network connections from the ESX server to each of the SAN servers. This redundancy will ensure that any single network connection—or even an entire SAN server—can be lost without affecting your running VMs.

Exiting this example's VMware ESX server are two NICs, both of which have been dedicated to iSCSI traffic. These two NICs aren't bonded through the use of traditional network teaming; iSCSI doesn't use traditional network teaming to aggregate its network connections. Instead, these two NICs will be bonded using iSCSI multipathing.

Unlike traditional NIC teaming, which presents only a single IP address to the outside world, iSCSI multipathing uses individual IP addresses for each connection—at both the initiator and the target. Thus, a server that enjoys a multipathed connection to an iSCSI SAN will need multiple assigned IP addresses. Each IP address will connect to an IP address on the SAN storage.

It's worth noting that although the SAN storage in this article's example uses only one IP address, it's doing so only for the sake of simplicity. Your production SAN storage should be configured with multiple IP addresses for redundancy, load balancing, and connectivity to multiple storage processors.

The following steps assume that you've completed VMware ESX's initial installation and that the server has been networked appropriately so that it can be managed through the vSphere Client.

1. Your first task is to configure the network cards that will be used for iSCSI storage. To do so, you'll use the vSphere Client. Click the Networking link on the Configuration tab. Then, click Add Networking and create a new VMkernel connection type. On the next screen, create a new virtual switch using only one of the NICs that you've identified for storage traffic.

2. On the next screen, you can label the port group with a friendly name. In Figure 3, you can see that the example's virtual switch is named iSCSI 1. If your environment uses virtual switch tagging to trunk VLANs to the VMware ESX server, you should also identify the correct VLAN ID in the box. Be aware that your network will need to be properly configured for VLANs to function. The VMware article "Sample configuration of virtual switch VLAN tagging (VST Mode) and ESX" (kb .vmware.com/kb/1004074) outlines the steps required to accomplish this task with your networking equipment.

3. Figure 4 shows the wizard's final screen, on which you'll need to configure your IP settings. Enter the IP address, subnet mask, and (optionally) the VMkernel

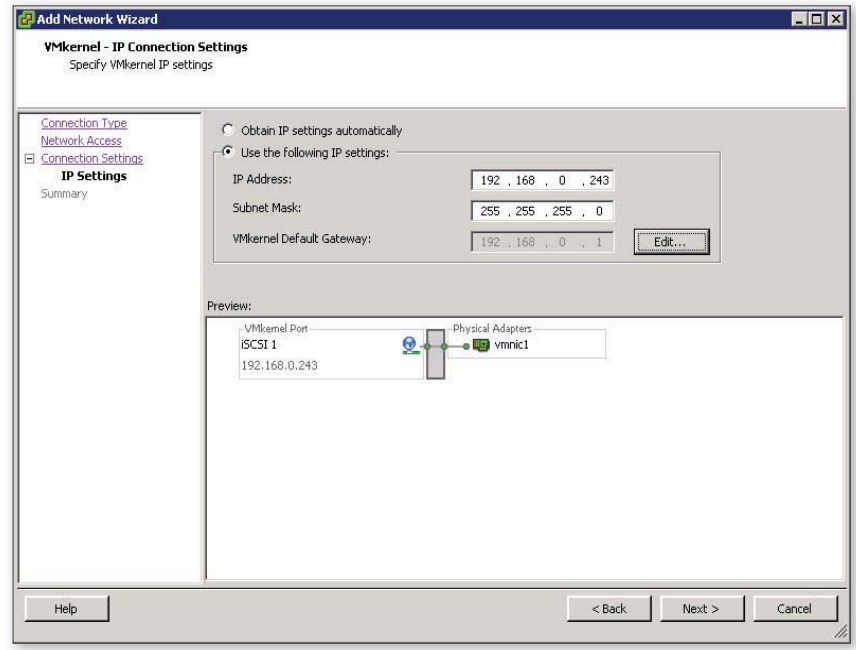

Figure 4: Configuring IP connection settings

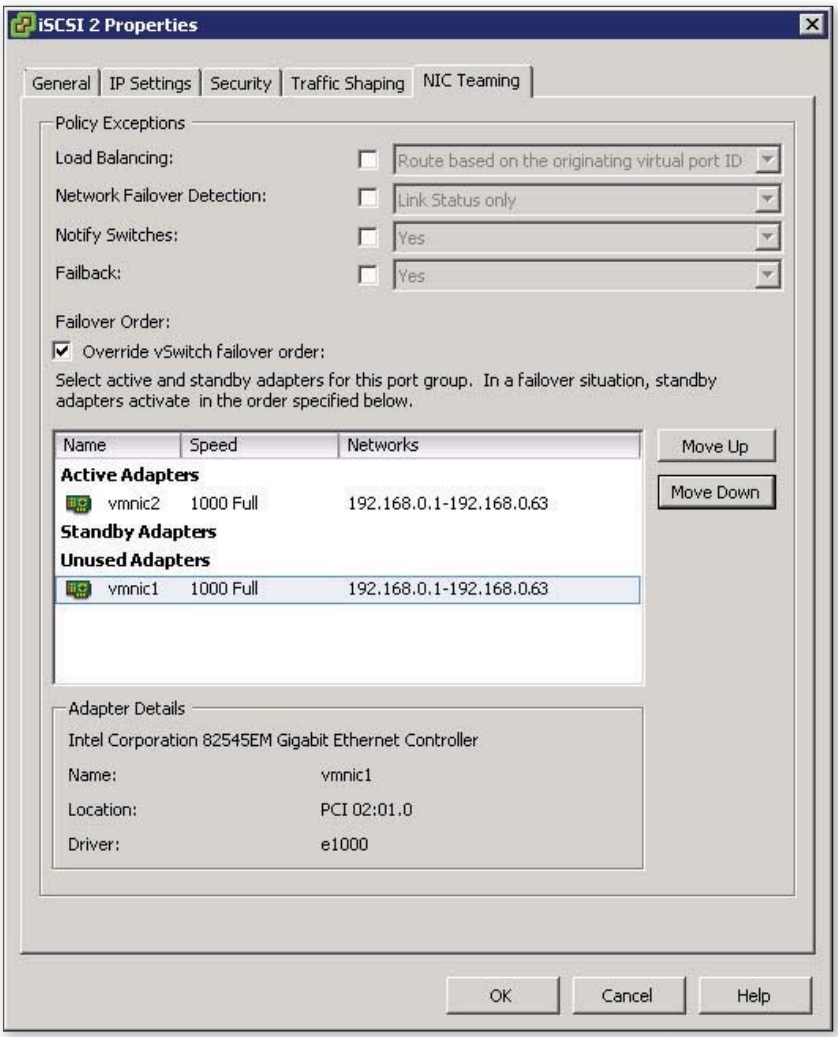

Figure 5: Overriding the vSwitch failover order

Default Gateway for the storage connection. Completing this step configures the first NIC.

4. To add the second NIC, access the properties of the Virtual Switch and select the Network Adapters tab. Click Add to add each additional NIC, then choose Next and Finish.

5. You must create a VMkernel port for each subsequent NIC. To do so, click the Ports tab in the Virtual Switch Properties console and select Add. Choose VMkernel for the connection type and enter the appropriate network label and IP address information for each subsequent NIC.

6. You need to create a 1:1 mapping between the NICs and VMkernel ports. By default, all network adapters will appear as active for each VMkernel port on the Virtual Switch. This doesn't work with iSCSI; iSCSI multipathing requires that you override this default setup so that each port maps to only one corresponding NIC. View the properties of the Virtual Switch again, and select the Ports tab. Select one of the VMkernel ports you just created (labeled in this article's example as iSCSI 1 and iSCSI 2), click Edit, and access the NIC Teaming tab. There, select the *Override vSwitch failover order* check box and ensure that only one NIC is set as an Active Adapter. Figure 5 shows how vmnic2 is set as the only active adapter for the iSCSI 2 port. Repeat this step for each NIC, ensuring that each NIC maps to only one port. Figure 6 shows the Virtual Switch configuration for this example's connection.

7. Now, you need to connect the VMkernel ports you just created to the iSCSI initiator. Start by enabling the iSCSI initiator itself. On the vSphere Client's Configuration tab, click Storage Adapters. Scroll through the list of storage adapters to find the iSCSI Software Adapter. Select this adapter, and click Properties, then Configure. Select the Enabled box, click OK, then click Close to enable the iSCSI initiator. Back at the vSphere Client, if you access the adapter's properties again, you'll see that it's now populated with a name and target discovery methods.

8. You'll need to use the vSphere command-line interface (CLI) to create the connection between the VMkernel ports you created in the earlier step and the iSCSI initiator. You can do so by logging on to the Service Console directly as *root*. The command syntax to accomplish this task is:

```
esxcli swiscsi nic add -n <port_name>
   -d <vmhba>
```

9. In Figure 6, you'll see that the two created ports were labeled vmk0 and vmk1. Now, take another look at the Storage Adapters screen in the vSphere Client. For example, if the iSCSI Software Adapter is set to vmhba33, you would use the following syntax inside the vSphere CLI. The first two commands make the connection, and the third command lists the results:

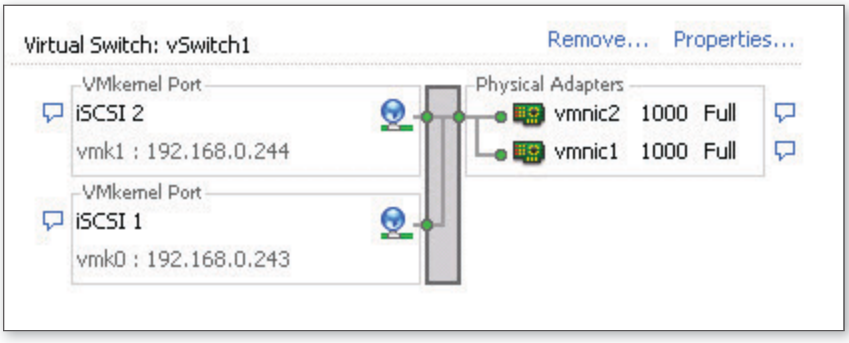

Figure 6: A fully configured virtual switch for redundant iSCSI traffic

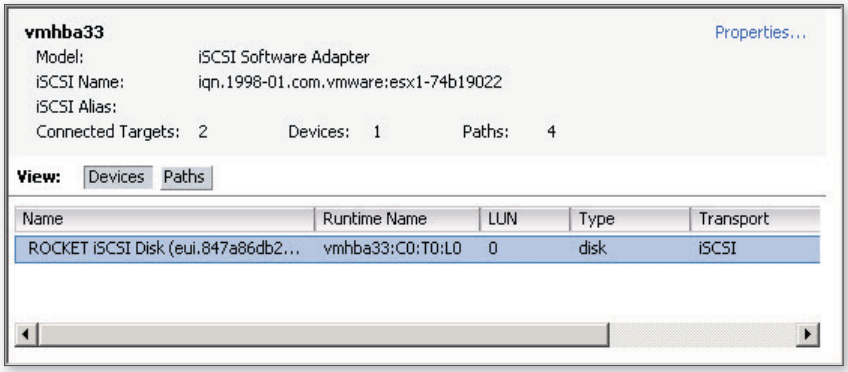

Figure 7: A configured LUN with four paths

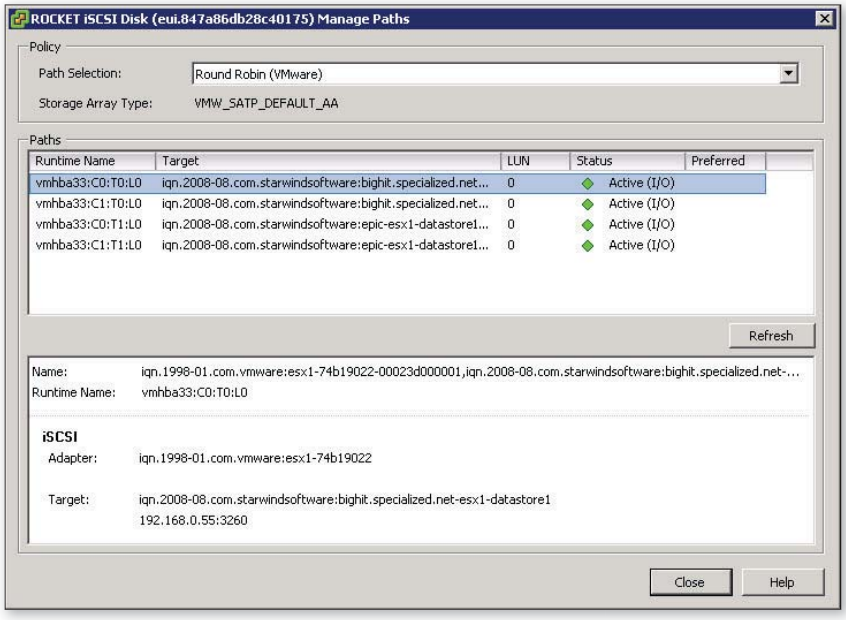

Figure 8: Changing the path selection for a LUN

```
esxcli swiscsi nic add -n vmkØ
  -d vmhba33
esxcli swiscsi nic add -n vmk1
  -d vmhba33
esxcli swiscsi nic list -d vmhba33
```

Connecting NICs to SAN LUNs

Your NICs are now ready for connecting to your iSCSI LUN. Recall that two servers comprise the StarWind Software SAN. Both servers will need to be addressed to create the highly available connection.

1. In the Properties console of the iSCSI initiator, click the Dynamic Discovery tab, then click Add. Enter the IP address for the StarWind Software SAN's iSCSI connection. (This will be the address you set in Figure 2.) Repeat this process for the partner SAN server.

2. Configuring a Send Target Server instructs the iSCSI initiator to send a Send Targets request to that server—essentially with the question *What LUNs do you have for me?* The server responds to that request by returning a list of available iSCSI targets that have been configured for the initiator. The initiator's Static Discovery tab displays the two targets that were sent back.

3. Click Close. The vSphere Client should present a dialog box prompting you to rescan the adapter to complete the configuration change. Choose Yes to rescan the adapter. If you've done everything correctly, the Storage Adapters screen should show a single LUN available to the VMware ESX server, as you see in Figure 7. Notice that four paths are available to the LUN. Those four paths correspond to the two storage NICs on the VMware ESX server, each of which is now connected to the two StarWind SAN servers.

4. You can right-click the LUN and select Manage Paths to perform additional configuration on the paths themselves, as Figure 8 shows. By default, an iSCSI connection will use the Fixed (VMware) path selection. This path selection instructs the host to always use the preferred path to the disk when that path is available, falling back to the other paths only when the preferred path goes down. The Fixed (VMware) path selection, as you can imagine, doesn't perform load balancing across your configured paths. To get load balancing, set the path selection to Round Robin (VMware).

5. Now, you need to add the newly connected storage to the VMware ESX server and create a datastore. To do so, click the Storage link on the Configuration tab in the vSphere Client. Then, click Add Storage and select a Storage Type of Disk/LUN. Note that not all SANs have special configurations over and above what I discuss in this article. However, some do. If you're using StarWind Software's solution, there are two more settings you'll want to add in the vSphere Client. Click the Advanced Settings link in the Software section on the Configuration tab. Configure the following two settings:

```
Disk.UseDeviceReset = Ø
Disk.UseLunReset = 1
```

6. Select the disk and complete the steps in the Add Storage wizard to create a new datastore on the LUN.

Redundant Redundancies

You're now ready to install VMs and enjoy your brand-new VMware ESX server! You have created a multiply redundant connection that can survive the loss of any network connection, or even a complete SAN server failure. You'll want to incorporate these same levels of redundancy into your production environment as well.

InstantDoc ID 125689

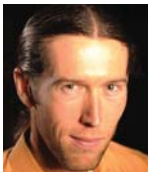

Greg Shields
(virtualgreg@concentratedtech.com) is an independent author, speaker, and IT consultant. He's a Microsoft MVP and a partner and principal technologist with Concentrated Technology.

Liberating Desktop Virtualization

Quest® vWorkspace. Master virtual desktop and application delivery through a single user access point and management console. vWorkspace blends Terminal Server/Remote Desktop Session Host, VDI, Blade/Physical PCs, and Application Virtualization into one solution. With the added freedom to choose from multiple virtualization platforms, vWorkspace delivers simplicity through consolidation. Get more with less complexity, resources and cost.

Experience how Quest liberates Desktop Virtualization Management. Read the white paper "Concept Becomes Reality With Quest" or watch the video at **quest.com/Liberating**.

QUEST SOFTWARE®
Simplicity At Work™

FEATURE

Are Your SSL Certificates Secure?

Current recommendations are to upgrade to 2,048-bit keys on all certificates

by Alan Sugano

If you review your network environment, you'll probably find that your company has several commercially issued SSL certificates and some self-signed SSL certificates. SSL certificates are used to secure things such as Microsoft Exchange Server's Outlook Web Access (OWA), SharePoint, SSL VPN appliances, and, of course, other websites. The National Institute of Standards and Technology (NIST) has issued a statement that says SSL certificates with a key length of 1,024 bits or fewer will be insufficient for security after December 31, 2010, because NIST estimates that computers will be powerful enough to perform a brute-force crack of keys of that size. For more information about this recommendation, refer to "Recommendation for Key Management" (csrc.nist.gov/publications/nistpubs/800-57/sp800-57-Part1-revised2_Mar08-2007.pdf). Surprisingly, there are quite a few large companies doing business on the web that still use SSL certificates with 1,024-bit keys.

You might be asking yourself, *Should I care?* Of course, only you can answer that question. But if you went through the hassle of installing an SSL certificate on a site in the first place, you probably do care about the security of your data on that site. So let's look at how you determine what the key length of your current certificates is and investigate some considerations you might need to address when updating 1,024-bit keys to 2,048-bit keys in your environment.

Determining SSL Key Length

You might be unsure of what key lengths your current keys have. One easy way to determine the key length of any SSL certificate is through Internet Explorer (IE) by following these steps:

1. Using IE, navigate to the site where the SSL certificate is installed.
2. Click the padlock (Security Report) immediately to the right of the URL, then click View Certificates.
3. Click the Details tab and scroll down until you see the *Public key* field. As Figure 1 shows, the SSL key length is shown in parentheses.

In this example, the SSL certificate was issued with a 2,048-bit key, so it complies with the NIST recommendation. In case you were wondering, NIST estimates that an SSL certificate with a 2,048-bit key will be viable until 2030. Of course, you can use this method to check not only your company's SSL certificates but also the SSL certificates of any company that has a secure website (i.e., a website that uses HTTPS). If you or your end users frequently visit secure websites (e.g., a 401k site), you might want to check the certificates for those sites to see if they comply with NIST's recommendations. If you find any that don't, you might consider contacting the companies to see when they plan to upgrade their SSL certificates.

Don't wait to reissue any 1,024-bit SSL certificates you find in your network because you could run into unforeseen problems that will delay the process. SSL certificates with a 1,024-bit key will

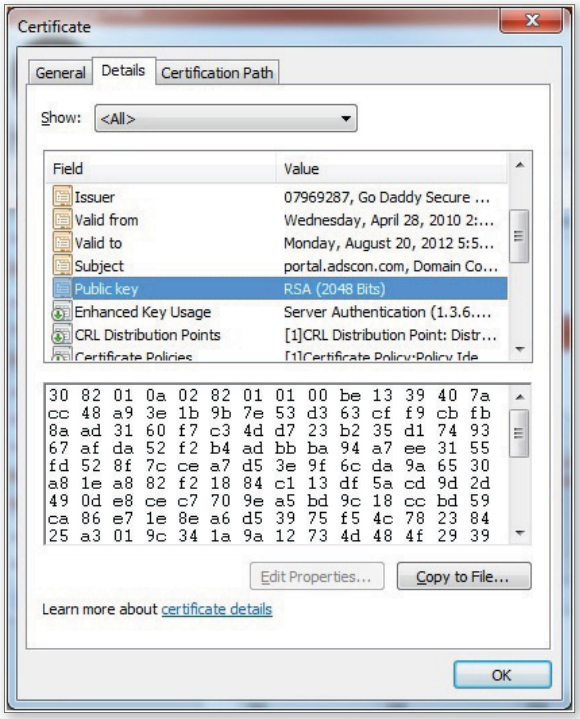

Figure 1: Checking the key length of an SSL certificate through IE

probably be more common for certificates that were renewed for a period of two years or more. For commercial SSL certificates, some vendors, such as Go Daddy, let you re-key an existing SSL certificate so that you don't have to purchase a new certificate if you just want to increase the key length of the certificate.

Load and Hardware Considerations

If your site has heavy SSL traffic, you might need a hardware upgrade before you increase the key length of the SSL certificate. Changing the key length from 1,024 bits to 2,048 bits places additional CPU load on the server or SSL appliance. In most cases, you'll probably be fine, but if your current hardware is barely keeping up with the existing SSL traffic, you might finally be able to justify the cost of the SSL accelerator appliance that's been shot down in your company's IT budget for the past two years. High CPU utilization on a web server might indicate a high SSL load—or it could possibly indicate something such as an errant script running on the web server. You can use Performance Monitor (perfmon.exe) with the Web Service counter to get an idea of the SSL load on the server.

A quick and dirty way to determine if your Microsoft IIS Server is experiencing high CPU load related specifically to SSL traffic is to download the Web Capacity Analysis Tools (WCAT) available from Microsoft at support.microsoft.com/kb/231282 and technet.microsoft.com/magazine/2008.04.utilityspotlight.aspx. Run the stress tools with secure and nonsecure sites with Performance Monitor active and measure the CPU load on the web server. If the CPU load stays close to 100 percent with HTTPS access and is significantly lower with the equivalent HTTP access, your load problems are probably related to your SSL traffic on the server. You can even run the stress tools with a 1,024-bit versus a 2,048-bit SSL certificate to see how the 2,048-bit SSL certificate will potentially impact a web server's performance.

In addition to the increased load to manage SSL sessions, you might discover that your current hardware doesn't support an SSL certificate with a key length longer than 1,024 bits. This limitation is completely independent of the current SSL traffic load on the device. For instance, one device that won't support a certificate with a key length of 2,048 bits is the SonicWALL SSL-VPN 200. According to SonicWALL, the problem is with the hardware, not the firmware, so there's no upgrade path for this appliance—you must purchase a new appliance that supports certificates that have key lengths of 2,048 bits. For more information about this problem, refer to www.fuzeqna.com/sonicwallkb/consumer/kbdetail.asp?kbid=7354. There are probably other appliances that fall into this category. It's not too early to start reviewing your SSL certificates and hardware so that you have time to budget, make purchases, and upgrade any hardware before the end of the year.

Generating the CSR

Most SSL vendors have stopped issuing SSL certificates with 1,024-bit keys and now require a certificate signing request (CSR) with a 2,048-bit key or longer. There are some SSL vendors (e.g., Thawte, VeriSign) that will still accept a CSR for a 1,024-bit key, but the 1,024-bit certificates such vendors issue will have an expiration date of December 31, 2010. When renewing your SSL certificates, it doesn't really make sense to generate a CSR with a 1,024-bit key length because the SSL certificate will be good only until the end of the year. If you need more information about SSL vendors, check out WhichSSL (www.whichssl.com) for good comparative information about different SSL vendors, which can help with your evaluation and selection.

If you have an existing SSL certificate on IIS that must be renewed, you typically generate a CSR renewal request and submit it to your SSL vendor. After the vendor validates your domain, the SSL vendor issues the SSL certificate. Different vendors have different processes for performing this validation. Most vendors send a message to the email address listed in the WHOIS field of the domain registration information for quick-issued SSL certificates; enterprise SSL certificates often have a much more stringent validation process. However, if you generate a renewal request for an existing SSL certificate that has a 1,024-bit key in IIS, the renewal CSR will also have a 1,024-bit key. The workaround is to export the existing SSL certificate, then generate a new CSR, which lets you select the key length.

Certificates with a key length of 1,024 bits or fewer will be insufficient for security after December 31, 2010, because NIST estimates that computers will be powerful enough to perform a brute-force crack of keys of that size.

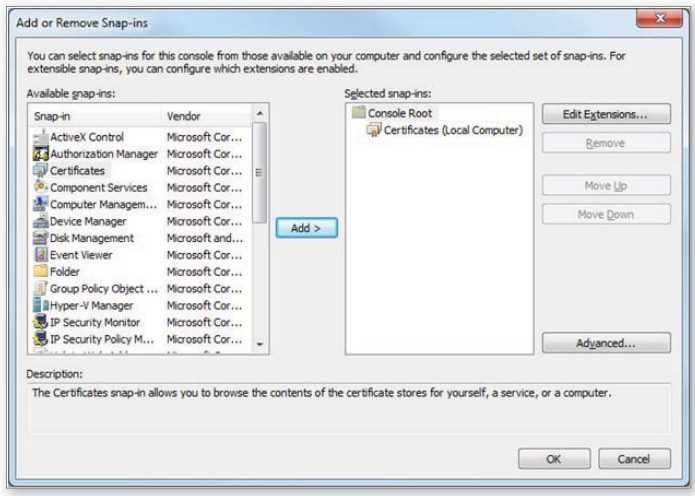

Figure 2: Adding the Certificates snap-in to MMC

Of course, your site will be down while the CSR is processed by your SSL vendor. For sites such as OWA, being down usually isn't a problem as long as the CSR is generated after business hours. Even for sites with heavy traffic, I still suggest generating a new CSR because I've found that it works reliably. If it's not possible to take down the site for any reason, you might need to configure a mirrored website, generate the new CSR, install the SSL certificate, test it, and then perform a cutover to the mirrored site to reduce the amount of downtime on the site.

For more information about this issue, check out the Microsoft article "How To Renew or Create New Certificate Signing Request While Another Certificate Is Currently Installed" (support.microsoft.com/kb/295281). For sites that can't be down for any length of time, it's especially important to be familiar with your SSL vendor's domain and company validation procedures. Make sure you have all the necessary email addresses and other validation mechanisms in place before you generate the CSR so the new SSL certificate can be issued on a timely basis.

Upgrading Root Certificates

Updates for root certificates can be downloaded from Microsoft Support (support.microsoft.com/kb/931125). Of course, you can also download root certificates directly from your SSL vendor. As you know, root certificates are used in the chain of trust to verify that your SSL certificate is valid. Because of this chain, you want to verify that the root certificate from your SSL vendor also has a key length of 2,048 bits or more. To check the key length of a root certificate, complete the following steps:

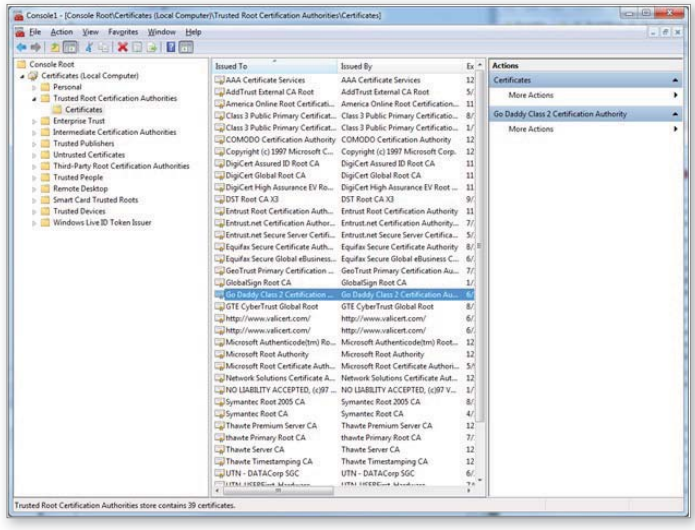

Figure 3: Using the MMC Certificates snap-in to review root certificates

1. Launch Microsoft Management Console (MMC).
2. In MMC, click File, Add/Remove Snap-in.
3. In the left column under *Available snap-ins*, click Certificates, then click Add.
4. Select *Computer account*, then click Next.
5. Select *Local computer*, then click Finish.
6. Figure 2 shows what the *Add or Remove Snap-ins* dialog box should look like. Click OK.
7. In the left pane of MMC, expand Certificates (Local Computer).
8. Expand Trusted Root Certification Authorities.
9. Select Certificates and you should see a list of root certificates, as Figure 3 shows.
10. In the center pane, double-click the root certificate you want to check. In this example, I checked the Go Daddy Class 2 Certification Authority.
11. Click the Details tab.
12. Scroll down to the *Public key* field and review the value.

In the Go Daddy root certificate I checked, the certificate had a key length of 2,048 bits, so the computer had the upgraded root certificate installed.

Get Ahead of the Curve for Security

It's not too early to start reviewing your SSL certificates to determine how many certificates you need to upgrade to the longer key length. What seems like a relatively simple process can get complicated if you have several SSL sites with heavy traffic that require hardware upgrades in order to be compliant with the NIST recommendations. However, I hope you're ahead of the curve in the process of upgrading your SSL certificates. Happy SSL upgrading!

InstantDoc ID 125820

Alan Sugano
(asugano@adscon.com) is the president of ADS Consulting Group, which specializes in networking, custom programming, Microsoft .NET web development, and SQL Server development. He's the author of *The Real-World Network Troubleshooting Manual* (Charles River Media).

Exchange 2010 MRM

Implementing New Retention Policies

How to design, create, and apply retention tags and policies

by Tony Redmond

We live in a world of ceaseless information. Our mailboxes are stuffed with messages, and we have little inclination to spend the time and energy to impose any order in our mailboxes. We keep everything until we're forced to free up some mailbox space when our quotas are exceeded. Outlook's automatic archive feature helps by regularly moving items from a mailbox into a PST, but it's a crude instrument because it moves items based on date rather than permitting more granular selection criteria. It also depends on the deployment of a single client, and that's just not a viable strategy in a world where client choice is the norm.

Microsoft Exchange Server 2007 introduced the messaging records management (MRM) system as its business email strategy to help users comply with regulatory and legal requirements. The concept is simple: automate the retention process so that users have an easy way to control what's in their mailboxes and retain those messages and attachments that are required business records.

To implement this concept, Exchange 2007 MRM uses a set of managed folders that have policies attached to them. The Managed Folder Assistant (MFA) applies the policies attached to the folders. Marked items are kept as long as required; other items are automatically discarded when their retention period expires. (For more information about Exchange 2007's MRM system, see the web-exclusive article "Using MRM to Manage Mailboxes" at www.windowsitpro.com/go/UsingMRM.)

Although the concept is good, in reality Exchange 2007's MRM system is typically ineffective because Exchange users are relatively undisciplined when it comes to filing messages in folders. So, Microsoft changed its tactics to provide a workable implementation of MRM in Exchange 2010.

A New Approach

Instead of using managed folders, Exchange 2010 MRM uses retention tags and policies. (Managed folders persist in Exchange 2010 but only for backward compatibility.)

Retention tags. Exchange 2010 supports three types of retention tags: retention policy tags (RPTs), default policy tags (DPTs), and personal tags. Table 1 describes these tag types.

Retention tags can be applied to any item in any folder to specify what action Exchange should take for the item when its retention period expires. Supported actions include the hard (permanent) or soft (recoverable) deletion of items, moving items to a personal archive, or flagging items for user attention. Retention tags can be placed on items, conversations, or complete folders. They're transferred with items if you move them between folders.

Retention policies. Retention policies group retention tags together so that administrators can apply policies to mailboxes rather than having to assign individual retention tags to folders. Retention tags and policies are organizationwide objects that are stored in Active Directory (AD) and can therefore be applied to any mailbox in the organization after they're created. The MFA is still responsible for checking mailbox contents against the policy and taking the specified action for each item that exceeds its retention period.

Retention polices and tags follow a couple of simple principles. First, Exchange can apply only one retention policy to a mailbox. If you apply a retention policy to a mailbox that already has one, Exchange

Table 1: Types of Retention Tags		
Tag Type	Use	Possible -Type Values
Retention policy tags (RPTs)	Administrators can apply these tags to default mailbox folders such as the Inbox, Sent Items, and Deleted Items. They can't be applied to the Calendar and Tasks folders. If an RPT is assigned to a default folder, all items in the folder automatically come under the control of the tag unless the user applies a personal tag to the item. Only one RPT can be assigned per default folder.	DeletedItems, Drafts, Inbox. JunkMail, Journal, Notes, Outbox, SentItems, All
Default policy tags (DPTs)	A catch-all tag that the Managed Folder Assistant (MFA) applies to any item that doesn't inherit a tag from its parent folder or hasn't had a tag explicitly applied to it by the user. In other words, if no other tag applies to an item, Exchange will respect the instructions contained in the DPT. A retention policy includes only one DPT.	All
Personal tags	Users can apply these tags to nondefault folders and individual items in a mailbox. Personal tags mark an item with an explicit retention, usually to comply with a business requirement. For example, you might use an "Audit" tag to mark items that users are compelled to retain for audit purposes. A retention policy can include many different personal tags.	Personal

will overwrite the policy already in place. Second, Exchange can apply only one retention tag to an item. If an item has more than one associated retention tag, Exchange uses two rules to determine which tag to apply:

1. The retention tag with the longest retention period always takes precedence. This rule ensures that Exchange never deletes an item before its time truly expires.
2. A retention tag applied specifically to an item (or all the items in a folder) takes precedence over one that an item inherits through a default retention policy. So, for example, if you apply a personal tag that states an item must be retained for 6 years but the default retention policy for the folder requires deletion after 12 months, the item will be kept for 6 years.

To use Exchange 2010 MRM, you need to deploy clients that include the necessary intelligence and UI. At the time of this writing, the only clients in this category are Outlook 2010 and Outlook Web App (OWA) 2010. Outlook 2010's UI provides the richest views of retention policies and tags. OWA 2010 is less capable but still very usable.

Designing a Retention Policy

Many different retention policy tags can exist within an organization, which allows great flexibility in creating appropriate policies for different groups that work within a company. For example, the finance department might want Exchange to permanently delete everything in the Deleted Items folder more than three days old (i.e., the shred principle), whereas other departments might not be concerned if items survive in the Deleted Items folder for 30 days or more. For the finance department members, you can apply a retention policy that includes an RPT that instructs the MFA to permanently remove items from the Deleted Items folder after three days. The same policy might include a personal tag that lets financial department members mark items that have to be archived for audit purposes after a month in their primary mailbox. The MFA will move items with this tag to the archive mailbox when it processes the mailbox.

Before you create a retention policy, you should determine the why, when, and how for the policy:

- Why are you implementing the policy? What business need will the policy serve?
- When will you implement the policy? What mailboxes will the policy be applied to? How will you communicate the policy to users so that they understand the purpose of the policy and how it will affect the contents of their mailboxes?
- How will you implement the policy? What tags are required? What actions will be enforced through tags and what retention periods will be used? Do any restrictions exist as a result of other aspects of your deployment? For example, if you use an archiving product from another vendor, you can't use tags to move items into an archive mailbox after a designated period.

The design for a retention policy can be captured in a simple table format that makes it clear what tags are included in the policy, their purpose, and the folders to which they apply. Apart from anything else, capturing the design like this makes it easier to communicate the policy to users. Table 2 lays out a sample policy that could be applied to help managers cope with their overloaded mailboxes. Logically, you should have only one RPT for each folder. It would be very confusing to have two RPTs compete within a single folder.

Although you can create and use as many retention policies as you want, the question of long-term supportability arises. A couple of well-designed, logical policies that satisfy the vast majority of requirements will be easier to manage than a mass of granular policies generated to meet the specific needs of each department or other business group. The more policies that exist, the more potential there is to confuse administrators and users alike.

Creating Retention Tags

Exchange 2010 MRM is sound in terms of both concept and implementation. The only problem is that Microsoft didn't ship a GUI that you can use to create and apply retention tags and policies in the release to manufacturing (RTM) version of Exchange 2010. Microsoft will provide the missing GUI in Exchange 2010 SP1. A beta version of this code was released at TechEd in June 2010, and the final version is expected to ship later this year. Until SP1 comes out, you can use PowerShell to create and apply retention tags and policies.

Let's create the various tags necessary to build the retention policy in Table 2. To create a retention tag, you use the New-RetentionPolicyTag cmdlet. So, for example, to create the RPT for the Inbox, you'd run the command

```
New-RetentionPolicyTag
  -Name 'Manager-Inbox'
  -RetentionAction MoveToDeletedItems
  -AgeLimitForRetention 30
  -Type Inbox
  -Comment 'Inbox items are
    automatically moved to
    Deleted Items after 30 days'
  -RetentionEnabled $True
```

Table 2: Sample Retention Policy

Retention Policy Name		Management retention policy	
Applies To		Mailboxes with CustomAttribute7 = "Management"	
General Purpose		Automatic clean-out of Inbox and Sent Items folders to encourage users to keep these folders tidy. Items in all other folders can remain in place for a year. Removal of items from the Deleted Items folder after a week and permanent removal of anything filed into the Junk Mail folder after two days. A tag is provided to allow users to mark items for retention for 5 years.	
Tag Name	**Tag Type**	**Applies To**	**Action**
Manager-Inbox	RPT	Inbox	Move items to Deleted Items after 30 days.
Manager-SentItems	RPT	Sent Items	Move items to Deleted Items after 30 days.
Manager-JunkMail	RPT	Junk Mail	Permanently remove items after 2 days.
Manager-Deleted	RPT	Deleted Items	Remove and allow recovery after 7 days.
Manager-General	DPT	All folders	Move items to Deleted Items after 365 days.
Manager-Retain	PER	All folders	Move items to Deleted Items after 1,825 days (5 years).

(Although this command wraps here, you'd enter it all on one line in the PowerShell console. The same holds true for the other commands that wrap.) The -Name and -Type parameters specify the tag's name and scope, respectively. The "Possible -Type Values" column in Table 1 lists the other possible values for the -Type parameter.

The -RetentionAction parameter specifies the action to take for items past the retention limit. In this example, the action is MoveToDeletedItems, which tells the MFA to move the item to the Deleted Items folder. The other actions you can specify are:

- MarkAsPastRetentionLimit. The MFA marks the item as being past its retention limit but takes no further action. Outlook indicates this status by striking a line through the item when it's included in a list of items.
- DeleteAndAllowRecovery. The MFA moves the item into the Deleted Items folder, but the user can recover the item with the Recover Deleted Items option if desired.
- PermanentlyDelete. The MFA immediately deletes the item in such a way that it can't be recovered using the Recover Deleted Items option. However, if the mailbox is on retention or litigation hold, the item is retained and still available to discovery searches.
- MoveToArchive. The MFA moves the item to a folder with the same name in an archive mailbox. This action is possible only when the mailbox has a personal archive. If not, the MFA ignores the action. The MoveToArchive action is similar to Outlook's AutoArchive option, which moves items into a PST on a regular schedule to help keep a mailbox under quota. However, unlike the AutoArchive option, the MoveToArchive action doesn't let uses decide whether they want to use that capability. The MFA automatically moves items into the personal archive without asking. Policies that move items into an archive mailbox are known as *archive policies*.

The -AgeLimitForRetention parameter specifies the retention limit. Exchange uses the date and time when an item (even a modifiable item such as a post) is created as the baseline to calculate the item's age for retention purposes. So, in the case of the Manager-Inbox RPT, the age limit of 30 days means that items are moved to the Deleted Items folder 30 days after they're delivered to the Inbox.

You can create a tag that tells the MFA never to process an item. To do so, you don't set a value for the -AgeLimitForRetention parameter and you set the -RetentionEnabled parameter to $False. If you set a value for the -AgeLimitForRetention parameter, you must always set the -RetentionEnabled parameter to $True.

The MFA ignores items that are older than the retention period when it processes a mailbox. For example, if the retention period for the Inbox is 30 days, the MFA will tag any item aged up to 30 days, take action for items aged 30 days, and ignore any item older than this. This approach might seem strange at first, but it's actually quite logical because you can't have a minus retention period. The upshot is the potential for user confusion because all the most recent items in a folder (up to the retention period) will be tagged while anything past this point will be ignored. Thus, when you implement a retention policy, you do so from a particular point rather than going back to the start of time.

After creating a tag, you can check its properties with the Get-RetentionPolicyTag cmdlet. For example, to the check the Manager-Inbox tag, you'd run the command

```
Get-RetentionPolicyTag
  -id 'Manager-Inbox'
```

Unlike managed folders, retention tags don't accommodate the notion of item segregation. In other words, you can't build a retention tag that only applies to items of a certain class in a folder (e.g., apply the policy to items of class IPM.Note, but ignore those of class IPM.Contact). Along the same lines, you can't define different actions for different item types, such as moving expired messages to an archive folder while deleting any other type of item. Some observers consider these shortcomings to be a step backward for MRM.

Using the New-RetentionPolicyTag cmdlet, you can create the remaining five tags listed in Table 2. You don't need to do anything different when creating a personal tag (e.g., the Manager-Retain tag) or DPT (e.g., the Manager-General tag).

To check that you have all the required tags in place to build the retention policy, you can execute the command

```
Get-RetentionPolicyTag |
  Format-Table Name, Type,
  RetentionAction, RetentionEnabled,
  AgeLimitForRetention -AutoSize
```

Creating the Retention Policy

Now that you've created the six tags required to help managers impose order on their mailboxes, you can create a new retention policy using the New-RetentionPolicy cmdlet. When you create the policy, you also associate its tags by using the -RetentionPolicyTagLinks parameters, as in

```
New-RetentionPolicy
  -Name 'Management retention policy'
  -RetentionPolicyTagLinks
  'Manager-Inbox', 'Manager-SentItems',
  'Manager-Deleted', 'Manager-JunkMail'
  'Manager-General', 'Manager-Retain'
```

You can use the Get-RetentionPolicy cmdlet to examine details of the new retention policy. For this example, you'd run the command

```
Get-RetentionPolicy
  -id 'Management retention policy'
```

A six-tag policy is a reasonably simple retention policy. Microsoft recommends that you include no more than 10 retention tags in a policy to avoid confusing users and administrators, which is good advice. There might be times when you need to incorporate more tags in a policy to meet specialized business needs. A more sophisticated policy for a department might have separate RPTs for all the default folders, a set of personal tags developed specifically to suit the department's retention needs, and a DPT for everything else.

Applying a Retention Policy

The Set-Mailbox cmdlet is used to apply a retention policy to mailboxes. A mailbox can have only one retention policy, so when you assign a retention policy to a mailbox, the action overwrites any policy that might already be in place. The new policy will be applied to the mailbox the next time the MFA runs.

The command to apply a retention policy might look like

```
Set-Mailbox -id 'JSmith'
  -RetentionPolicy
  'Management retention policy'
```

Exchange will warn you that clients earlier than Outlook 2007 don't support retention policies. No client before Outlook 2010 and OWA 2010 includes the UI necessary to display the retention policy information about items and change the retention tag that's applied to an item or folder.

If you're setting a policy for a group of users, you can do it in one operation by selecting the mailboxes with the Get-Mailbox cmdlet and piping the results to Set-Mailbox, as in

```
Get-Mailbox -Filter
  {CustomAttribute7 -eq 'Management'} |
Set-Mailbox -RetentionPolicy
  'Management retention policy'
```

This operation will become a lot easier in Exchange 2010 SP1, when the GUI will be available. You'll be able to select a group of mailboxes and apply a retention policy using Exchange Management Console. EMC will also let you apply a retention policy to an individual mailbox.

To discover the set of mailboxes that have retention policies in place, you can use a command like this:

```
Get-Mailbox -Filter
  {RetentionPolicy -ne $Null} |
  Format-Table Name, RetentionPolicy
  -AutoSize
```

After you begin to deploy retention policies to mailboxes, you need to determine how to integrate the assignment of retention policies with your company's user provisioning process. Exchange doesn't have a default retention policy that can be assigned automatically, so an explicit administrative action is always required to assign a retention policy to a mailbox. As you've just seen, this task isn't difficult to code with PowerShell, but it's something that needs to be considered as part of your deployment plan.

Upgrading from Managed Folders

You can upgrade a managed folder to a retention tag by using the folder as the template for the new tag. The New-RetentionPolicyTag cmdlet's -ManagedFolderToUpgrade parameter makes this possible. Suppose you have a managed folder named Never Delete that acts as a repository for items that users never want to be removed from a mailbox because they're so important. You could argue the case that these items could be stored in an archive mailbox. However, archive mailboxes didn't exist in Exchange 2007, and it takes time for people to change their behavior. A better alternative might be to create a new retention policy tag based on the Never Delete managed folder by using the command

```
New-RetentionPolicyTag
  -Name 'Mark item to never expire'
  -ManagedFolderToUpgrade 'Never Delete'
  -Comment 'Tag created from old Never
  Delete managed folder'
```

As you can see, when you use the Managed FolderToUpgrade parameter, you don't need to set the -Type, -RetentionAction, -AgeLimitForRetention, and -Retention Enabled parameters. This information will be retrieved from the specified managed folder.

To complete the process, you have to associate the new tag with a retention policy and apply the retention policy to the users' mailboxes. After this is done, the users will be able to apply the new tag on any item in their mailboxes.

Stay Tuned for More

The topic of Exchange 2010's retention policies is too big to be covered in just one article. So, I'll cover how to change and remove retention policies, as well as how to help users understand them, in the article "Exchange 2010 MRM: How to Modify and Reduce Help Desk Calls About Retention Policies," which will appear in the November issue of *Window IT Pro.*

InstantDoc ID 125359

Tony Redmond
(12knocksinna@gmail.com) is a contributing editor for *Windows IT Pro* and the author of *Microsoft Exchange Server 2010 SP1 Inside Out* (Microsoft Press, October 2010).

Debugging in Windows PowerShell

How to squash existing bugs and prevent new infestations

by Don Jones

If you ever tried your hand at writing a PowerShell script, you likely had to spend some time debugging it. Bugs are an inevitable part of life when you're trying to tell a computer something to do. On the surface, PowerShell doesn't seem to offer much in the way of debugging assistance. In this short guide, I'll tell you about some of the basic techniques for debugging and some practices that can help keep the bugs at bay.

Although PowerShell doesn't have any kind of one-line-at-a-time graphical debugger built in, there are free and commercial third-party tools that do. Editors such as Quest Software's PowerGUI (powergui.org), Idera's PowerShellPlus (powershellplus.com), and SAPIEN Technologies' PrimalScript (www.primaltools.com) use different techniques to trick PowerShell into running your script one line at a time, showing you what a script's variables contain and generally making debugging easier. However, unless you know what debugging is all about, these tools don't make you a more effective debugger.

Before you start debugging, you need to know what you're looking for—in other words, what is a bug? There are two distinct types.

Bug Type One: Fat Fingers

The first type of bug is a simple syntax error. A typo. An operator malfunction. Fat fingers. Pick a term. You'll find these bugs relatively easy to spot because PowerShell will hurl an error in your direction, complete with the number of the line that contains the error. However, if you're using Notepad, knowing the line number won't help because it doesn't display line numbers. For this reason, you should stop using Notepad and use PowerGUI, which is free, or a commercial script editor instead.

Commercial script editors offer more than just line numbering, and some of their features can help prevent fat-finger bugs. Chief among these capabilities is syntax highlighting, which is nothing fancier than making your script code turn colors. But the key is that things only turn the right color when they're spelled correctly. So, if you're using a syntax-highlighting editor, start becoming familiar with the colors it uses for command names, variables, literal strings, and so forth. If they aren't turning the right color as you type, you missed something. Go back and look carefully for the problem.

If you do miss something, PowerShell will be more than willing to tell you about it. Pay attention to error messages. I can't tell you how often I see admins struggling to fix something, simply because they're not reading the error message. When they see that red text hit the screen, they go into a bit of a panic and just start trying different things. But the error message is telling you what line of code has the problem, and most of the time it's even telling you what the problem is. So just slow down a bit, read carefully, comprehend what the error is saying, and see if you can spot the problem.

A good script editor can provide further help by asking PowerShell to run a sort of preflight checklist on your script. This feature is typically called live syntax checking. The editor can then call

your attention to simple errors right away, before you even run the script. Some use a red underline (like Microsoft Word's spell-check feature); others use different visual indicators.

Finally, adopt some best practices to keep yourself out of trouble. Format your code nicely so that constructs are indented and their curly braces line up, as in

```
Function Get-Something
{
  # Code here
}
```

rather than messy code like this:

```
Function Get-Something { # Code here }
```

The first example is easier to read and easier to verify that there's an opening and closing brace. When you start nesting constructs, it's easy to miss a closing brace when code is set up like the second example.

Bug Type Two: Expectations Don't Match Reality

The second type of bug has to do with expectations not being met. With these bugs, your script runs without complaint, but it doesn't do what you think it should. PowerShell typically doesn't return an error message with these types of bugs, or it returns an error message that doesn't make sense. There are only two reasons why this type of bug exists:

You made a logic error. People sometimes make basic logic errors. For example, let's say you want your script to count from 1 to 10, so you run the code

```
For ($i=0; $i -lt 10; $i++)
{
  Write-Host $i
}
```

All goes well when the script runs, except that the results show 0 to 9. Close, but not quite what you wanted. This is a case where slowing down and thinking like the shell can pay off. "Thinking like the shell" means pretending you're PowerShell while reading the script, taking the time to write down what the variables should contain at each step as well as each step's results. For the script just given, the run-down might go as follows:

I first need to initialize $i to 0 (write down *$i=0*). Is 0 less than 10? Yes, so I'll use Write-Host to display the value of 0 (write down *result is 0*) and increment $i by one. At this point, $i contains 1 (write down *$i=1*). Is 1 less than 10? Yes, so I'll display the value of 1 (write down *result is 1*) and increment $i by one. Now, $i contains 2 (write down *$i=2*), which is less than 10, so I'll display that number (write down *result is 2*), and so on.

This would continue until $i contains 10. At this point, you'll find that 10 is not less than 10, which is why PowerShell didn't display that number. As this example shows, most basic logic errors become obvious when you take the time to do this simple run-through.

> To find basic logic errors, you should "think like the shell"—pretend you're PowerShell while reading the script, taking the time to write down what the variables should contain at each step and each step's results.

You made an assumption. Thinking like the shell can solve not only basic logic errors but also more serious logic errors because the paper record provides a list of your expectations. Let's say you've run through your script in your head and written down what you think should happen, but when you run your script, something else happens. I refer to this kind of bug as a "bad expectation" or "lack of knowledge" bug. The cause is simple: A variable, property value, or method result doesn't contain what you thought it would contain. Because you've already done the act of "thinking like the shell," you have a written record of your expectations. So, all you need to do is have Power Shell generate the same kind of list, then see where its results differ from your expectations. When you find the difference, you'll have found the bug.

There are two basic techniques for getting PowerShell to generate the same kind of list you created by thinking like the shell. The techniques involve using the Write-Debug and Set-PSDebug cmdlets. There's actually a third technique, which involves using a script editor's interactive debugger. That experience is basically just a roll-up of the two underlying techniques I'm about to cover—and I want you to know the underlying techniques before you start taking the easier road with an editor's interactive debugger.

Technique 1: Use Write-Debug

The first technique involves using the Write-Debug cmdlet. By default, output from this cmdlet is suppressed. To enable that output, you need to add the line

```
$DebugPreference = "Continue"
```

to the top of your script.

Then, you need to start adding Write-Debug statements immediately after any statements that change a variable's contents, as in

```
$var = Get-Service
Write-Debug "`$var contains $var"
```

Notice the little trick I used: I enclosed the Write-Debug cmdlet's output in double quotes, which tells PowerShell to replace the variables with their actual contents.

Listing 1: Adding Write-Debug to *for* and *if* Statements

```
for ($i=0; $i -lt 10; $i++)
{
A  Write-Debug "`$i is $i"
   if ($i -gt 2) {
     # Do something here.
B    Write-Debug '$i is greater than 2'
   } else {
     Write-Debug '$i is not greater than 2'
   }
}
```

For the first variable, though, I used the backtick (PowerShell's escape character) so that the *$* was escaped. This will result in the first $var being displayed as-is, so you can see the variable name. The second $var will be replaced with its contents.

Next, you need to add a Write-Debug statement in every loop and decision construct. Listing 1 shows examples of how to add it to a *for* loop and *if* construct. In the Write-Debug statement in callout A, I use the same double-quotes-and-backtick trick to find out what $i contains in each iteration of the *for* loop. In the *if* construct, I use two Write-Debug statements in an *else* construct, as callout B shows. With this setup, I get debug output no matter which way the decision goes, even though I don't have any code to execute if $i is not greater than 2. Notice that I use single quotes for the Write-Debug statements' output so that the $i variable wouldn't be replaced with its contents.

After all the Write-Debug statements have been added, it's time to run the script and compare the debug output with your written expectations. Look for any differences. Perhaps you'll need to revise your expectations, but any differences are a likely place for a bug to live.

When your script is running as expected, simply change $DebugPreference at the top of your script to

```
$DebugPreference = "SilentlyContinue"
```

and your debug output will be suppressed again. There's no need to remove the Write-Debug statements. Leaving them in will make debugging the script easier in the future.

Technique 2: Use Set-PSDebug

The second technique involves using PowerShell's built-in step debugger, Set-PSDebug, to go through your script one line at a time. This can be a bit tedious in long scripts, but in PowerShell 1.0, it's what you have to work with. (PowerShell 2.0 also supports breakpoints, which I'll be covering in a future article. They allow you to define when your script will enter suspend mode, rather than forcing you to step through it one line at a time.)

To enable the debugger, run the command

```
Set-PSDebug -step
```

You can disable it later by running

```
Set-PSDebug -off
```

Once enabled, the debugger is effective for all scripts. (It also works on commands entered on the command line.) You just run your script and start debugging. Each time PowerShell encounters a line in your script, you'll see that line displayed and PowerShell will ask if you want to continue. (I like having a line-numbered printout of my script handy so that I can see exactly where PowerShell is.) To run a line, press Enter. This is where you can start comparing your expectations with what PowerShell actually does.

Any time your script changes a variable or is about to use a variable or object property, don't press Enter. Instead, type S then press Enter. This suspends your script and presents you with a special prompt. In it, you can:

- Access variables to see what they contain
- Run commands to see what they produce
- View object properties to see what they contain

When you find a difference between the shell's results and your expectations, you've found a bug. Run the Exit command to leave the suspend mode and resume script execution right where you left off.

The step debugger takes a bit of time to get used to. However, it can be useful to get inside your script while it's actually running.

Ready to Debug?

Debugging can seem painful, but it doesn't need to be, provided you're methodical, consistent, and patient in your approach. When you receive an error message, don't let the red text alarm you. Take the time to read and understand it. Treat it as a friend that just wants to help (albeit in a somewhat garish, annoying manner).

When you don't get the results you expect, take the time to document your expectations and find where your expectations don't match reality. The number one cause of bugs is someone who has no idea what to expect from the script at each step of the way. That's common for newcomers who are trying to, for example, use a script they found on the Internet. By taking the time to understand a script, you'll not only be helping yourself debug faster but also helping yourself learn more about PowerShell and become more effective at writing your own scripts in the future. In that regard, debugging can be an investment that's well worth the payoff.

When you receive an error message, don't let the red text alarm you. Take the time to read and understand it. Treat it as a friend that just wants to help (albeit in a somewhat garish, annoying manner).

InstantDoc ID 125694

Want to ask Don Jones a question about PowerShell debugging? Submit your question at windowsitpro.com/go/SubmitFAQ.

Don Jones
(powershell@concentratedtech.com) is the author of more than 35 books and is a speaker at technology conferences such as Microsoft TechEd and Windows Connections. He's a multiple-year recipient of Microsoft's MVP and is technical guide for PowerShell at www.windowsitpro.com/go/DonJonesPowerShell.

FEATURE

Exchange Server's Client Access: Load-Balancing Your Servers

With all client connections coming through this server role, it's crucial to ensure that your servers can handle the load

by Ken St. Cyr

The Client Access server role plays a big part in Microsoft Exchange Server 2010 by providing the access point for every Exchange client. With such a big responsibility, you need to ensure that your Client Access servers can handle the load from your users and that these servers have minimal downtime.

In my previous articles in this series, I provided an introduction to the Client Access server role ("Exchange Server's Client Access: An Introduction," www.windowsitpro.com/article/exchange-server/Exchange-Server-s-Client-Access-An-Introduction.aspx) and gave you an overview of how to deploy it in your organization ("Exchange Server's Client Access: Deploying Your Servers," www.windowsitpro.com/article/messaging/Exchange-Server-s-Client-Access-Deploying-Your-Servers.aspx). In this article, I focus on giving your Client Access servers a little R&R—that's Resiliency and Redundancy. I'll do this by showing you how load balancing works and then how to apply it to the Client Access role.

Introduction to Load Balancing

At a high level, load balancing is a way to distribute a workload across multiple systems. Sometimes people refer to these systems as a *farm* or an *array*. By distributing the load, you maximize your servers' resource use while minimizing response delay and system downtime. In a load-balanced array, redundancy is acquired inherently: If one of the systems in the array goes offline, that system's load transfers automatically to another system.

Network-based load balancing is a form of clustering, but you shouldn't confuse it with the concept of resource clustering. The difference between the two is that in resource clustering there are resources that the groups of servers share. These resources can take many forms, such as system services, IP addresses, or data. In network-based load balancing, however, there are no shared resources across the nodes. Instead, load balancing is TCP/IP-based. A user accesses a network service on a specific port, such as an HTTPS connection on port 443, and the load balancer ensures that one of the servers in the array responds to that connection.

Load balancers come in two primary forms: external devices and server software. Different load-balancing products not only work differently but also scale differently and provide unique features. In general, though, the premise is the same—they intercept network traffic before it reaches the service and distribute the load across the array of servers.

External load balancers sit in front of the server array; instead of accessing the individual servers directly, the end user or reverse proxy accesses the load balancer. Figure 1 shows a load-balanced array of Client Access servers with an external load balancer. In this design, the load balancer has a unique name (mail.contoso.com) that users access directly. When a connection is established with the load balancer, the device determines which Client Access server will do the work, then brokers the connection to that server node.

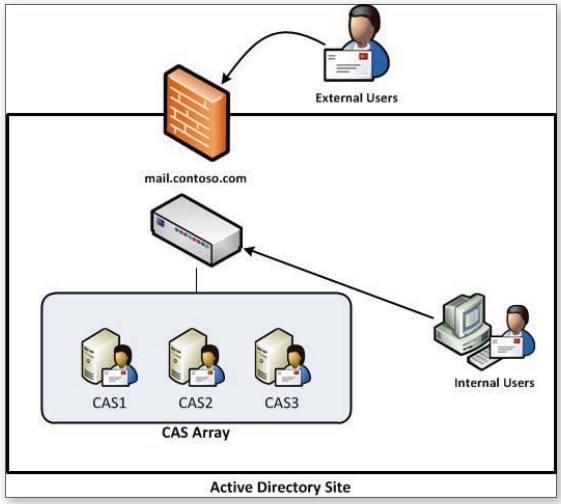

Figure 1: A load-balanced array of Client Access servers with an external load balancer

External load balancers have different methods to determine which Client Access server to direct the connection to. For example, they can use predefined rules, such as client subnet mapping, where one server is responsible for all traffic from a specific range of IP addresses, or they can take a round-robin approach. Different types of load balancers can perform different types of monitoring, such as determining if a service is running or if a port is accessible on the Client Access server, and some products can use such monitoring to determine which server to access.

A software-based load balancer is installed on each of the servers in the array. One common software-based load balancer you might use in smaller Client Access server implementations is the Network Load Balancing (NLB) service that's built in to Windows Server. Software load balancers typically use predefined algorithms to determine which server should handle user requests. Figure 2 demonstrates a typical Windows NLB configuration. In this case, the user connects to a virtual name and IP address that's shared by every server in the array. This shared IP address is assigned a unique MAC address that's used by every node, either in unicast or multicast mode. In unicast mode, the MAC address of one of the network adapters in each node is replaced with the shared MAC address. In multicast mode, there's an additional (multicast) MAC address added, so the network adapters retain their original MAC address as well. In either case, each server receives the network traffic for the shared IP address. To determine which node in the load-balanced array handles the network packets, a filtering algorithm is used. For Windows NLB, this algorithm is based on the IP address of the client that's connecting to the array. Every node in the array uses the same algorithm, so only one node will respond to the packet and the others drop it.

Windows NLB typically isn't suitable for most deployments. Although it technically supports as many as 32 nodes, you should never use Windows NLB to load-balance more than 8 Client Access servers because there could be scalability problems beyond that. Microsoft recommends you use a hardware load balancer if you have more than 8 Client Access servers in an array. Another disadvantage of Windows NLB is that it has limited built-in intelligence. It can monitor whether a port is up or down, but that's it. If a service crashes or something else breaks on the Client Access server, Windows NLB might still think the service is running and continue to direct clients to the server. Also, for Exchange deployments where you have a couple of all-in-one servers that participate in a Database Availability Group (DAG), you can't use Windows NLB because it's incompatible with Windows Failover Clustering.

Understanding Persistence

Load balancing works very well when no session-specific data needs to be maintained for each connection, such as when connecting to a farm of web servers that host static data. This situation doesn't apply to Client Access servers, however. Outlook Web App (OWA), Exchange Control Panel (ECP), and Exchange Web Services (EWS) require session state to be maintained. To understand how session-specific data affects these services, let's look at an example. Suppose there are two Client Access server nodes in an array (CAS1 and CAS2) and the virtual name of the array is mail.contoso.com. A user accesses https://mail.contoso.com/owa to log on to her web mail account. When she accesses mail.contoso.com, the connection is handled by CAS1. She enters her credentials and establishes an authenticated session. What would happen if the user opened a message and the load balancer decided that CAS2 was going to handle the request? Because session data isn't shared between CAS1 and CAS2, the user isn't authenticated with CAS2. Instead of the message opening, she would be prompted for re-authentication.

To solve this problem, load balancers employ a technique called *persistence* or *sticky connections*, which ensures that after the connection is established with CAS1, that connection always goes to CAS1 while the session is open. The only acceptable time the connection might be shifted to CAS2 is if CAS1 goes offline. In that case, the user would re-authenticate and continue working in CAS2. For an external load balancer, persistence is configured in the device itself and not on the Client Access server. Typically, external load balancers achieve persistence based on the client's IP address, the session ID of the SSL connection, or through the use of cookies, with cookie-based persistence the most common technique for load-balancing OWA.

In Windows NLB, persistence is configured in the affinity setting. You have three options when configuring affinity—Network, Single, or None. Table 1 outlines what each of these affinity settings does.

Load-Balancing HTTP Traffic

Load-balancing techniques and configurations differ between different Client

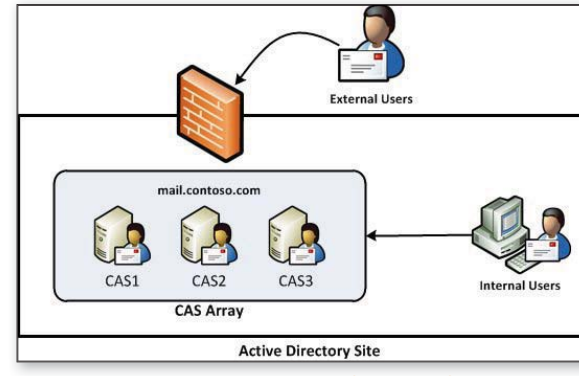

Figure 2: A typical Windows NLB configuration for Client Access servers

Access server services. Let's look at the HTTP-based services that should be load-balanced and discuss how to handle each one.

OWA and ECP. OWA and ECP connections can be load-balanced like a standard web application. You'll want to ensure that you use persistence when load-balancing these services so that session data is maintained. If you don't, users might be randomly asked to re-authenticate when they're directed to another server in the array. Cookie-based persistence is commonly used with external load balancers. Make sure you use the same load-balancer configuration for both OWA and ECP to ensure that a user who clicks the Options button in OWA doesn't get directed to another server when connecting to ECP. For both OWA and ECP, you want to load-balance TCP port 443.

Exchange Web Services. EWS connections can be load-balanced in a manner similar to OWA. If an EWS client makes a call that keeps some form of stateful data in memory on the Client Access server (such as a notification-based event, also known as a subscription), then there needs to be assurance that the client will hit that same Client Access node on subsequent calls. Some EWS clients won't require persistence, but some will. Therefore, you'll want to provide persistence on your load balancer for EWS. However, EWS clients might not be able to accept cookies, so you'll need to use persistence based on SSL session ID or client IP address. To ensure that EWS requests proxied between sites are persistent, EWS uses a special parameter on its virtual directory called InternalNLBBypassUrl. When you install the Client Access server, this parameter is set to the internal name of the Client Access server. This parameter is used to ensure that the correct Client Access node is used in EWS proxy scenarios, so you shouldn't change this URL.

Outlook Anywhere. If you used Outlook Anywhere in Exchange 2007 without persistence or with SSL session ID–based persistence, you could run into remote procedure call (RPC) problems around DSProxy. RPC connections are full-duplex connections: They require that data can be sent and received at the same time. HTTP doesn't allow for such transmissions because it's only half-duplex. So to simulate the required behavior in RPC over HTTP, two connections are established—RPC_IN_DATA for the incoming connection and RPC_OUT_DATA for the outgoing connection. Each of these connections is associated with a session ID for the clients. When the RPC component receives these connections with a matching session ID, it knows it needs to reply to RPC_IN_DATA requests over the RPC_OUT_DATA connection. If the RPC endpoint is the same for RPC_IN_DATA and RPC_OUT_DATA, it doesn't matter which Client Access server the connection is brokered through. Both the Information Store (port 6001) and the referral service (port 6002) had no problems with this in the past.

However, in Exchange 2007, DSProxy (port 6004) simply proxied these connections rather than being the actual endpoint. Because of this setup, the RPC_IN_DATA and RPC_OUT_DATA connections would sometimes be established with different domain controllers (DCs), breaking directory connections. There were some workarounds in Exchange 2007 to prevent this from happening, but the workarounds caused additional risks, such as tying Outlook profile creation to a single DC.

In Exchange 2010, this problem is resolved because DSProxy is no longer used. Instead, referrals are used for directory connections. The Name Service Provider Interface (NSPI) now exists on the Client Access server in the form of the Address Book service, so the client establishes the directory connection with the Client Access server and the Client Access server connects to a DC over LDAP. Even though persistence isn't required, you might still choose to use it to minimize session handshakes, although cookie-based persistence can't be used with Outlook Anywhere.

Table 1: Windows NLB Affinity Settings

Affinity Setting	Method Used for Persistence
Network	Persistence is established based on the subnet that the client is coming from.
Single	The client IP address is used to maintain persistence with the node in the array. If the client IP address changes during the session, the client might be directed to another node.
None	Persistence is guaranteed only within the connection. If a new connection is established and the source port is changed, the client can potentially be directed to another node in the array.

Load-Balancing MAPI Traffic

Having to load-balance MAPI traffic is a new problem we face in Exchange 2010. Until now, MAPI traffic went directly to Mailbox servers, so no form of load balancing was necessary. However, because MAPI clients connect to the Client Access server through the RPC Client Access service in Exchange 2010, you need to ensure that load balancing is in place if you want a fully redundant and highly available configuration. So there's a new Active Directory (AD) object called the Client Access array object introduced in Exchange 2010 for MAPI load balancing. This object lets you address a group of Client Access servers in a site as a single array. You can't have more than one Client Access array per AD site.

MAPI traffic uses RPC, which works differently from the HTTP traffic used by other Client Access services. RPC is used to execute code on another system on the network or in a different address space on the same system. To establish an RPC connection, a request is made on the RPC endpoint mapper port, 135. From there, a port is pulled from a dynamic range (ports 1024–65535) and that port is assigned to the RPC connection. Subsequent communications for this connection occur over the assigned port.

Because of this model, RPC requires a wide range of ports to be open to clients. Therefore, when load-balancing MAPI over RPC, all of these ports—135 and 1024 through 65535—need to be specified in the configuration. However, you also have the option of statically defining the RPC ports used. If you do this, the load balancer needs to be configured for both mailbox access and the Address Book service because they use separate RPC connections. To set a static port for mailbox access, you need to create a DWORD registry key called TCP/IP Port on

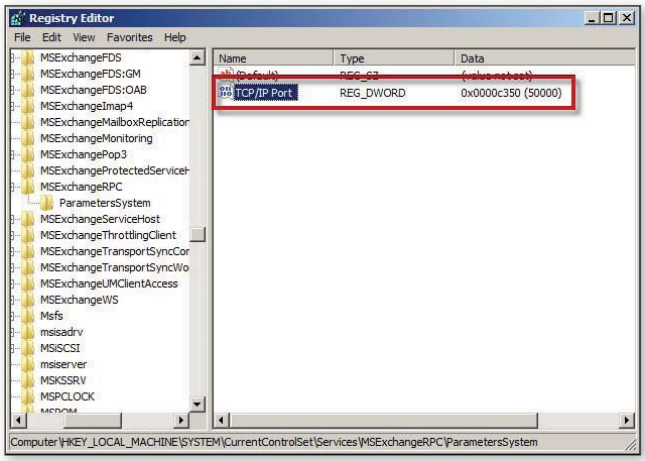

Figure 3: Setting a DWORD value to create a static port for mailbox access

your Client Access server in the following registry location: HKLM\System\CurrentControlSet\Services\MSExchangeRPC\ParametersSystem. You have to create the ParametersSystem subkey if it doesn't already exist. For the value of this key, enter the port number you want to use. As Figure 3 shows, I'm using port 50000.

For the Address Book service, you need to edit the XML file named microsoft.exchange.addressbook.service.exe.config in the Bin folder where you installed Exchange. For a default installation, that path would be C:\Program Files\Microsoft\Exchange Server\V14\Bin\. Open this file in Notepad and find the line that reads

```
<add key="RpcTcpPort" value="0" />
```

Change the value from 0 to the number of the port that you want to use. Reboot the Client Access server after making these changes.

If you set static ports, ensure that you make this same change on every Client Access server in the array. You can then configure your load balancer to load-balance only these ports instead of the entire range. When using static ports, you'll want to ensure that the load balancer is configured for port 135 in addition to the two static ports that you defined.

After the load balancer is configured, you're halfway to having your MAPI clients load-balanced. To complete this process, you need to create a Client Access array object that's assigned to the AD site that these Client Access servers exist in. You can do this with the following Exchange Management Shell (EMS) command:

```
New-ClientAccessArray
  -FQDN outlook.contoso.com
  -Site "Baltimore"
```

Next, you have to configure each of the mailbox databases in that site to use that Client Access array. This step is important because this parameter is used to tell your Outlook clients which Client Access server is used for directory connections—remember, the Client Access server now contains the NSPI endpoint. If you already have the Client Access array object created and assigned to a site, all new mailbox databases created in that site are automatically configured correctly. However, if you had databases in existence before you created the Client Access array object, you can configure the database to use the array with the following EMS command:

```
Set-MailboxDatabase DB01
  -RpcClientAccessServer outlook
     .contoso.com
```

Putting It All Together

Now that we've looked at how load balancing works and understand the factors that affect load-balancing Client Access servers, let's put all the pieces together in a series of steps. You can use the following process to load-balance your Client Access servers.

1. Install the Client Access servers and optionally configure static ports for RPC connections by MAPI clients.
2. Create the DNS entries for your load-balanced names. In this article, I used mail.contoso.com for web-based clients and outlook.contoso.com for MAPI clients.
3. Install and configure the load balancer. Table 2 summarizes the ports and persistence settings that you need for each Client Access service.
4. Create the Client Access array object using the New-ClientAccessArray cmdlet.
5. Configure your existing mailbox databases in the site to use the new Client Access array object.

When you understand the configuration nuances of each Client Access service, setting up a load-balanced Client Access server array isn't difficult. Vendors of hardware load balancers might have some specific guidance for configuring their particular device for Exchange 2010. Therefore, ensure that you understand your vendor's recommendations in addition to the information that I've provided in this article.

InstantDoc ID 125863

Table 2: Ports and Persistence Settings for Client Access Services

Client Access Service	Ports	Persistence Needed?
Exchange ActiveSync (EAS)	443	No
Outlook Anywhere	443	No
Outlook Web App (OWA)	443	Yes
Exchange Control Panel (ECP)	443	Yes
Exchange Web Services (EWS)	443	Yes
MAPI (dynamic)	135, 1024-65535	No
MAPI (static)	135, user-defined ports	No
POP3	110, 995 (SSL)	No
IMAP4	143, 993 (SSL)	No

Ken St. Cyr
(ken.stcyr@microsoft.com) is a solution architect at Microsoft with more than 10 years of industry experience. He's a Microsoft Certified Master in Directory Services and the author of *Exchange Server 2010 Administration Instant Reference* (Sybex).

WinConnections ...

Providing the **vision**

THE CONVERSATI

Microsoft® Exchange CONNECTIONS

UNIFIED COMMUNICATIONS CONNECTIONS

WINDOWS CONNECTIONS

SQL Server CONNECTIONS

SharePoint CONNECTIONS

NOVEMBER 1-4, 2010

LAS VEGAS • MANDALAY BAY RESORT & CASINO

QUESTIONS ANSWERED • STRATEGY DEFINED • RELATIONSHIPS BUILT

WINDOWS CONNECTIONS //

- Featuring the industry's most well-known and respected technology experts
- Learn to "do more with less" while increasing your IT skills
- Original content you won't find anywhere else
- Opportunities to make new connections amongst your peers
- Connect with key technology solution vendors

FIVE KEY FOCUS AREAS:

- Business Technology
- Building Your Skill Set
- Windows Server 2008 R2
- Windows 7
- Virtualization

Exchange Connections & Unified Communications //

Exchange and OCS Solutions for the Real World:

- Deployment
- Management
- Maintenence
- Microsoft Business Productivity Online Services (BPOS)
- Information Protection
- Useful Features in Service Pack 1
- Integration of Exchange with SharePoint (and other collaboration solutions)
- Best ways to use Unified Communications in your organization
- Sessions that Cover Exchange 2003 and Exchange 2007 and preparing for Exchange 2010
- Celebrate the release of Unified Communications 14

SharePoint Connections //

- Upgrade and Deployment to SharePoint 2010
- Enterprise & Web Content Management
- LINQ
- Business Connectivity Services
- Silverlight
- Workflow
- Virtualization
- Claims-based Authentication
- Enterprise Search
- Business Intelligence
- Security
- SharePoint Connections Bonus! No Code Solutions Track

intelligence
to keep you and your company
competitive in today's market!

ON BEGINS HERE

Only Microsoft and Industry Experts speak at WinConnections!

A sampling of our speakers ...

QUENTIN CLARK
MICROSOFT

STEVE FOX
MICROSOFT

DAN HOLME
INTELLIEM, INC.

KIMBERLY L. TRIPP
SQLSKILLS.COM

PAUL S. RANDAL
SQLSKILLS.COM

PAUL ROBICHAUX
TRAINER / AUTHOR

DON JONES
CONCENTRATED TECHNOLOGY

ITZIK BEN GAN
SOLID QUALITY LEARNING

RHONDA LAYFIELD
CONSULTANT/TRAINER

ALAN SUGANO
ADS CONSULTING GROUP

MARK MINASI
MR&D

TONY REDMOND
TONY REDMOND AND ASSOCIATES

JOEL OLESON
QUEST SOFTWARE

CHRIS AVIS
MICROSOFT

TODD O. KLINDT
SHAREPOINT911

ENJOY A PREMIERE LAS VEGAS HOTEL!

Mandalay Bay Resort And Casino

The Mandalay Bay Resort and Casino offers elegance, excitement, and escape. Indulge in award-winning restaurants helmed by celebrity chefs, an enormous beach-front pool, the luxurious spa, and top-notch entertainment. You'll find plenty of ways to relax with your colleagues from around the world at the end of each of our knowledge-packed days.

CHECK WEB SITE FOR DESCRIPTIONS OF SESSIONS AND WORKSHOPS

www.WinConnections.com • 800.505.1201 • 203.400.6121 • Register Today!

Microsoft · SharePoint Pro Connections · SQL Server Magazine · Windows IT Pro · Tech Conferences Inc. · Penton Media

FEATURE

Escaping SharePoint Permissions Purgatory

Don't let your environment become a place in which permissions run amok

by Ryan Thomas

If you use SharePoint 2007 extensively as a collaboration tool, you're probably having some trouble managing permissions. And if you upgraded your environment from Windows SharePoint Services (WSS) 2.0 or SharePoint Portal Server 2003, your challenges are likely even more extensive. Almost every organization must break permission inheritance at the site level to take advantage of SharePoint's security features. When breaking permission inheritance in SharePoint, you increase flexibility but often at the expense of maintainability.

My goal in this article is twofold. First, I want to shed some light on how the problem can manifest, while providing some advice about how to decrease the rate at which it grows. Second, I'd like to share the thought process that my company went through toward a solution that brought permissions back to a more manageable state, and I'd also like to provide some options for periodically reporting on the current status.

The problem I refer to is a lack of central management of permissions leading to a significant number of custom permission levels and permission assignment across a large number of sites. This can easily lead to inadvertently adding or removing access to sensitive data, site lock-out, content deletion, and duplication of efforts. A consistent management and naming convention for permissions combined with selective use of the Manage Permissions role can help halt these problems.

How Permissions Work

Although SharePoint 2007 offers extensive flexibility in the realm of permissions management, this flexibility can create a breeding ground for maintenance problems—particularly when site owners don't fully understand how permissions work and what potential damage can be caused. SharePoint uses several approaches to permissions management: base permissions, permission levels, permission assignment, inheritance, and item-level permissions.

Base permissions. At the root level, a series of base permissions dictate specific rights, and sets of these rights make up permission levels. These hard-coded, out-of-the box base permissions—which can't be added to—represent the building blocks of creating rights for users and groups. Figure 1, Figure 2, and Figure 3 show these permissions. SharePoint administrators have some control over the use of these base permissions. For example, the Full Control permission level is a common role that includes some potential dangers, such as the Manage Permissions base permission. This gives users the right to create their own custom permission levels.

Permission levels. Permission levels are groups of base permissions bundled together to assign to users or groups. More than one permission level can be assigned to a user, Active Directory (AD) group, or SharePoint group. Figure 4 shows a sample of permission levels, including a custom one I created. These should be familiar to most users.

Figure 1: Personal permissions

Permission assignment. When groups, users, or AD groups are assigned or used in SharePoint, they must be given a permission level within the site collection. You assign rights to content in SharePoint by assigning one or more permission levels to users and groups. Figure 5 shows an example, including a custom group with multiple permission levels assigned.

Inheritance. The previously discussed permissions information might seem straightforward, but two important permission-related factors remain. The first is inheritance. SharePoint wouldn't be very useful if all sites and content were forced to share the same permissions. By default, all sites and content inherit all the permissions from the site above it—all the way to the root web in the site collection. SharePoint provides mechanisms to break this inheritance at any or all of the sites and to add custom permission assignments. Generally, this capability might manifest itself as removing certain groups and adding new ones. These groups are often appropriately named for their level of access permissions relative to the site. If you have a large number of sites, this process can generate a significant number of SharePoint groups. If you choose not to use groups, you must assign rights to individual users or AD groups, which can create additional overhead.

Breaking inheritance often leads to the perceived need for additional permission levels, as well. For example, site owners might want certain users to be able to add items but not edit or delete them, but they might not want others to have access to version history. If site owners contain the Manage Permissions right in their permission level, they can create a custom permission level at will. One of the most important elements of this action is that permission levels created in sites with broken inheritance live only within that site. Therefore, site collection administrators have no idea how many custom permissions have been created, have been assigned, or have control, because the Manage Permissions right can give any user the same power to create groups and permissions.

Item-level permissions. If the previous options weren't enough to make you nervous, knowing that inheritance can be broken at the site, list, or list-item level for each and every site in each site collection in all web applications should raise the hairs on your neck. The upshot is that each item in your lists and document libraries has its permissions managed individually, including adding and assigning custom permission levels to custom groups.

Where's the Problem?

Because permissions management is so flexible, you'll often find a significant amount of unintended abuse of the permission assignments in SharePoint. In WSS 2.0 and SharePoint Portal Server 2003, the concept of a permission level didn't exist as it does in SharePoint 2007. So, SharePoint farms that went through upgrades containing unique permission assignments ran headlong into a complicated situation. Those permissions automatically rolled up into new permission levels during the upgrade—but the "permission level" name is generic, essentially displaying permissions with no useful name to identify intentions. These permission levels are also created at the individual subsite level, making them difficult to find.

Previous versions of SharePoint also allowed Microsoft Outlook integration, which could expand groups into individual users and assign them all to SharePoint. It's not unheard of to see 1,000 or more users all assigned with the same permission level to a site or list. Ideally, these users are added to a group to more easily identify and manage access to content.

One important question to ask is, "Do you know the status of your own SharePoint environment?" Of course, wrapped up in that question are related questions: How many users have the Manage Permissions right? How many lists or list items across your farm have unique

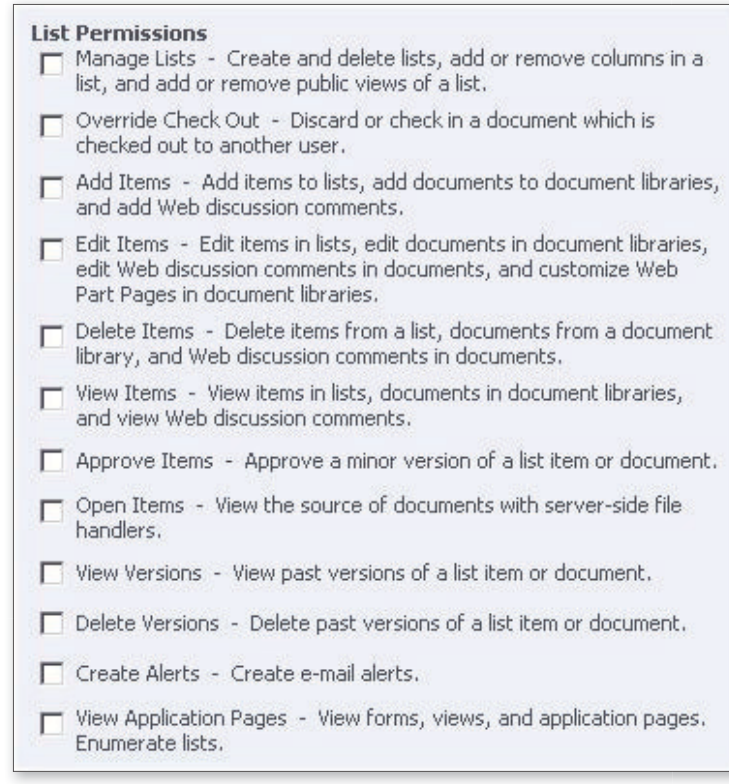

Figure 2: List permissions

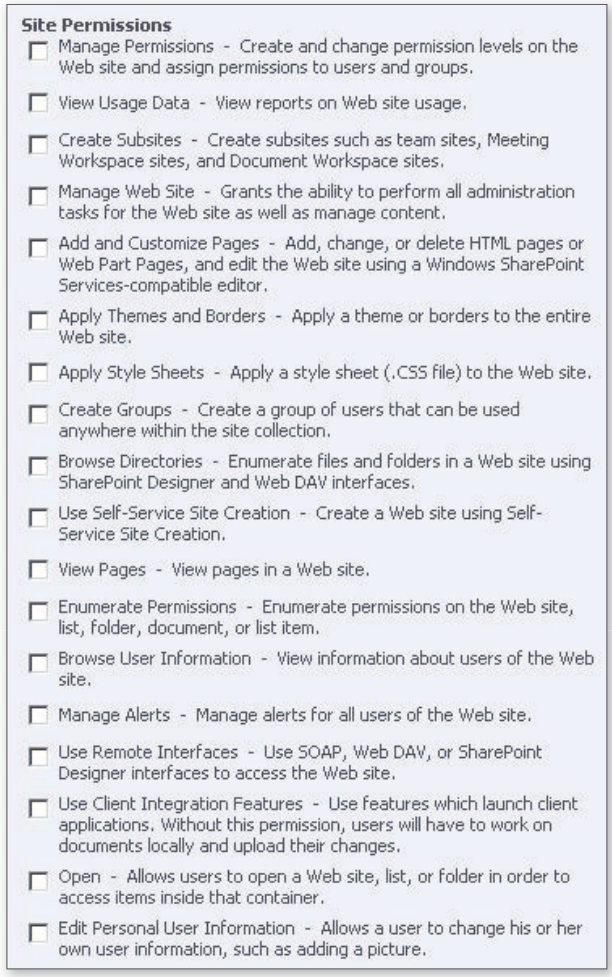

Figure 3: Site permissions

permissions? How many distinct SharePoint groups exist? How many locations have broken inheritance? At what level—just site, or list and list items? How many permission levels have been created without your knowledge with duplicate names but not duplicate rights, or duplicate rights but not duplicate names? How do you keep track of this?

Perhaps you still want to know why this is a problem. The main problem with permissions run amok is maintenance and security. Users accidentally delete, update, corrupt, move, or change content and sites. This scenario creates a burden on administrators tasked with protecting and managing the SharePoint environment.

The easy problems are those that are identified; the difficult problems are those that go unseen. For example, say a site owner gives inexperienced users additional rights to a team site, and that team site is accidentally deleted. The problem is identified and corrected after a certain amount of effort depending on the backup-and-restore strategy in place. But what happens if that user is accidentally deleting list items or documents? Those items might go unnoticed and cause unforeseen catastrophic issues with no visible explanation. From a security perspective, giving too many rights to manage permissions also leads to giving unintended access to content.

What's the Solution?

The problems I've described create varying levels of complications for different types of organizations. To assist with this problem in my environment, I ended up building a small application to identify the underlying data, analyze the permissions status, and create a system for fixing the permissions. Although third-party applications can assist with viewing, copying, and managing permissions, I resorted to writing my own application because nothing I found was robust enough for what I wanted to do.

The application isn't the focus of this article, but SharePoint development experience will be necessary should you choose to write your own. Ultimately, this was the only solution that met all the requirements to fix our problem—even after we researched commercial applications—but individual organizations will have different needs. Hopefully, some of the experience and insight I offer in this section will compel you to take measures now that prevent your organization from resorting to extreme measures to correct permissions problems. Now, here are the steps our team took.

1. Consolidate permission levels—Our first step was to take control of all the permission levels in our large SharePoint site collection. Our team had a vision of consolidating unique sets of rights down to six or eight that we could name, describe, and share with all the subsites. As I mentioned earlier, subsites that have decided to not only break permission inheritance but also edit permission levels cause an additional problem: Permission levels will no longer be shared across sites. Essentially, that means we can't enforce unique sets of permissions across sites. This problem is caused by allowing users to have the Manager Permissions base permission.

2. Consolidate into groups—Having consolidated permission levels to a manageable number of permission sets and forcing all subsites to use them, we wanted to eliminate the problem of assigning rights to individual users and assigning the same user multiple permissions. The objective was to have only SharePoint groups assigned to permissions and to add all users to these various groups. This would be a gigantic undertaking if performed manually.

3. Ensure proper access—One of the problems with consolidating users into groups and eliminating most of the existing permission levels is that we would either be granting additional permissions or taking away some rights from a number of users. This problem is inevitable, but there's a logical workaround for it. Our

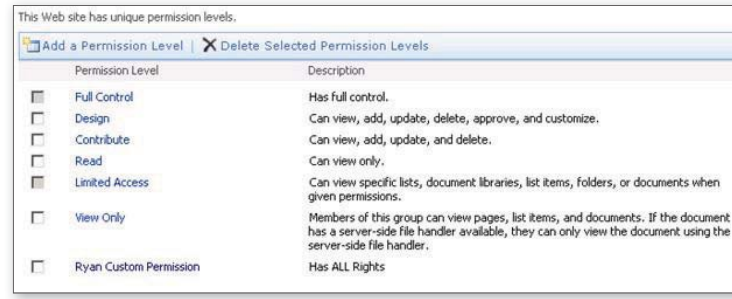

Figure 4: Permission levels

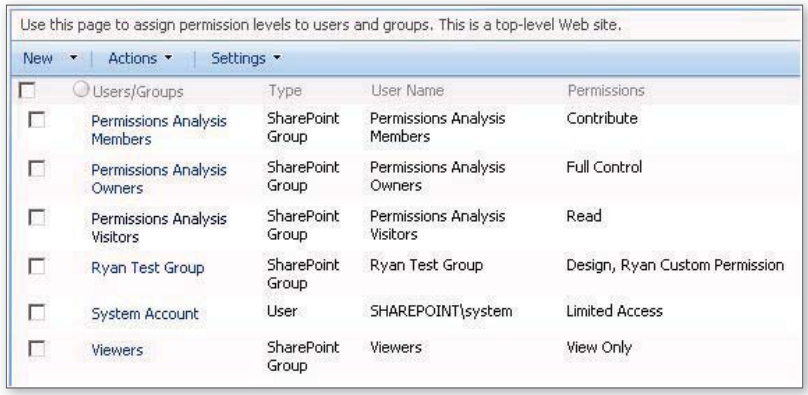

Figure 5: Permission assignment

Figure 6: Creating a matrix

approach was to use the data we gathered to see all the unique permissions assigned throughout the entire site collection and then break the permission assignments down into the base permissions. Doing so made it relatively easy to use a matrix to identify 10 to 15 permission levels that should actually be a single level. To help visualize the situation, Figure 6 shows a screenshot of the portion of a spreadsheet used to create this matrix.

As you can see, most of the permissions are similar and it became easy to bucket them into permission levels. We were able to look at 11 unique sets of rights and see that they were all intended to be a Read right. An added benefit of the Read right is consistency: We could make sure all the base permissions assigned to this level were standard across all sites.

After identifying all the unique rights, we were able to match them all down to a reasonable number. The only out-of-the-ordinary levels we had to account for were some additional rights for surveys to restrict the ability to delete or edit items, and some administration rights to restrict the ability to manage permissions and create subsites. By limiting the permission levels to specific rights with proper naming conventions, we also simplify the viewing of site permissions. For example, a SharePoint group might be assigned to the following levels: Site Owner, Manage Permissions, and Create Sub-site. We have been very careful granting Manage Permissions rights moving forward.

How Did We Do It?

Let's walk through the application so that you can understand how we accomplished all that. The application iterates through each site, list, and list item in the site collection and looks for broken inheritance. If it finds a broken permission, it stores the specifics of any role assignments for the given object in the database. Role assignments in the object model involve attaching a user or group to a permission level on a certain item (e.g., site, list, list item).

In this step, we also capture the specifics of the item, including permissions. So, in essence, we're cataloging all the places permissions are set and storing the data in a SQL Server database. Doing so allows us to build and run some relatively simple SQL queries to see all the unique permissions across the site collection. Once we've gathered the data, we can use it to build the matrix, then update the database to tell the application that each unique permission level should be mapped to the new consolidated level.

Now, we had the basis for where we wanted everything to end up. Plus, we had a catalog of what was stored in a database. The next step was to build the code that would implement all the desired changes across the site collection. This was the tricky part, because we had to reinherit a site structure in order to centrally manage these new permission levels and ensure that they were propagated to all sites. Therefore, we had to ensure that we had properly stored all our broken permissions so that they could be reapplied with the new permission levels.

As a part of this step, we had to write code to look at the role assignments to consolidate users into groups. Groups didn't exist in many cases, so we created them dynamically based on a naming convention of site name and permission level. Then, we added all users with that set of permissions to the new group. We did this after we ran the matrix to ensure that the new manageable permission levels had been created.

When we finished, we were able to meet the goals I outlined earlier: limiting certain unneeded permissions, rolling users into properly named groups, reassigning all existing permissions into the new matrix, and managing permission levels centrally.

Long-Term Considerations

As long as users have the Manage Permissions right, they will slowly be able to undo this work. I recommend limiting this right to only users who need it and understand the implications. One of the benefits of building a custom application is that you can periodically run the analysis portion to see how your work is holding up. The analysis will quickly identify the location of additional rogue permissions and any users assigned this permission—enough data to track down where problems are so that they can be fixed before the problems get too large to handle manually.

InstantDoc ID 125599

Ryan Thomas
(rthomas@syrinx.com) is director of the SharePoint Practice at Syrinx Consulting. He's a Microsoft Certified Professional Developer and Microsoft Certified Application Developer, and he contributes regularly to the Syrinx SharePoint blog and other industry publications.

- Virtualization
- Cloud Computing
- Archiving
- Security

3PAR Announces Support for VMware vSphere 4.1

3PAR has announced full support for VMware vSphere 4.1 and the development of a new **3PAR plug-in for VMware vStorage APIs for Array Integration** (VAAI). The 3PAR plug-in supports three features of VAAI: Hardware Assisted Locking, Block Zero, and Full Copy. In addition, the 3PAR plug-in supports greater scalability by preventing VMs from competing from the same resources, increased performance and efficiency by eliminating repetitive write commands, faster VM cloning and VMware Storage vMotion, and new quality of service capabilities via support for VMware Storage I/O Controls (SIOC). To learn more about the 3PAR plug-in, visit www.3par.com.

Centrify Corporation Releases Centrify Express

Centrify has announced **Centrify Express**, a set of software applications, tools, content resources, and community forums designed to help organizations improve security and compliance of data center and desktop systems. In this release, Centrify Express includes DirectControl Express, which lets Linux and Mac OS X servers and workstations participate in a Windows Active Directory domain. In addition, the release includes DirectManage Express, which automates deployment and management activities by providing a central solution to discover non-Windows systems on a network. Centrify Express also allows for cross-platform connectivity and file sharing with Centrify-enabled versions of OpenSSH, PuTTY, and Samba. For more information, visit www.centrify.com.

Red Gate Releases Exchange Server Archiver 3.0

Red Gate Software has released **Exchange Server Archiver 3.0**, a Microsoft Exchange email archiving software solution. The new release of Exchange Server Archiver is five times faster than its predecessor and requires no extras for the installation. In addition, Exchange Server Archiver can find and import PST files using the PST importer. Finally, Archiver 3.0 is scalable with the enterprise and supports Microsoft Exchange 2010. To learn more about Exchange Server Archiver 3.0, visit www.red-gate.com.

Symplified Announces New Capabilities for SinglePoint

Symplified recently announced several new capabilities for its **SinglePoint** identity and access management (IAM) solution. These new capabilities—Symplified Sync, Symplified Identity Vault, and the SinglePoint Virtual Directory—provide capabilities for managing and synchronizing user identities in on-premises IT infrastructures or cloud applications. In addition, the Symplified Identity Vault for Google and Salesforce.com transforms these two cloud applications into a cloud directory service for managing user accounts and serving as an authentication mechanism for other applications. For more information, visit www.symplified.com.

Dell KACE Offers Secure Browser

Dell KACE has released **Secure Browser**, a tool that protects against common security threats while browsing the web. The release of the Secure Browser includes an installation package that includes the Firefox browser, Adobe Reader, and Flash plug-ins that are pre-installed. Also, the browser provides easy website restriction policies that allow only specified sites to be accessed or restricted. In addition, the browser allows full remote insight and control when used with the K1000

PRODUCT SPOTLIGHT

Diskeeper Corporation Releases V-locity 2.0 with VMware Support

Diskeeper Corporation has released **V-locity 2.0**, a virtual platform disk optimizer that is designed to deliver invisible background optimization of all Windows Guest OSs running on VMware ESX and Microsoft Hyper-V platforms.

"Fragmentation clogged disk subsystems can lead to an inability to run more VMs on given hardware infrastructure, and lead to disk performance bottlenecks for VMs that share a common storage subsystem," notes Diskeeper Product Manager Michael Materie. "V-locity is designed to alleviate the 'virtual' disk bottleneck for VMs and provide a faster and more efficient computing platform for new consolidation and provisioning initiatives, without having to add more hardware."

New to V-locity 2.0 is the addition of IntelliWrite fragmentation prevention technology. V-locity utilizes this technology by writing files to the disk to prevent fragmentation from occurring. V-locity also frees up storage resources by eliminating virtual disk "bloat." V-locity does this by compacting the virtual disk, thereby preventing waste and allowing users to better allocate their virtual storage resources.

Also, Diskeeper Corporation's InvisiTasking technology allows background applications to operate with zero impact on a system. As more VMs are added to a host platform or dynamically migrated to new hosts, InvisiTasking will dynamically adjust to changing environments, allowing V-locity 2.0 users to optimize their virtual disk platforms.

For pricing and volume licensing discounts, call 800-829-6488, or visit www.diskeeper.com.

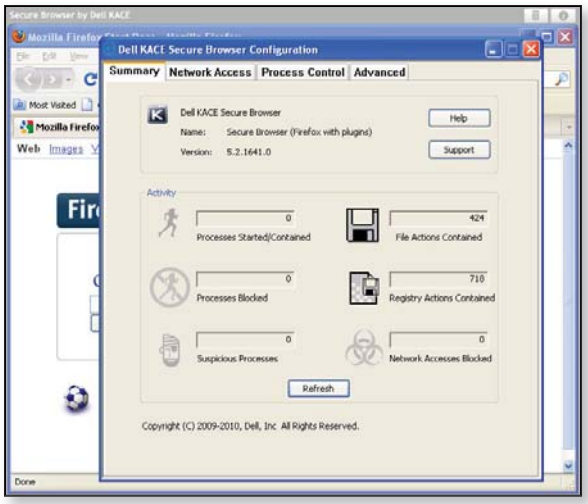

Management Appliance for advanced management capabilities. To learn more about the Dell KACE Secure Browser, visit www.kace.com.

REFOG Releases REFOG Terminal Monitor

REFOG announces **REFOG Terminal Monitor**, a centralized PC surveillance and employee tracking software for Terminal Servers. The REFOG Terminal Monitor helps administrators monitor employee productivity, prevent breaches of security, and enforce corporate policies by tracking, logging, and reporting all or select activities of all users logged onto the Terminal Server. In addition, the Terminal Monitor does this through application tracking, intercepting URLs of resources visited, monitoring chats, and intercepting keyboard strokes. To learn more, visit www.refog.com.

Comodo Announces Endpoint Security Manager Version 1.6

Comodo has announced the release of **Endpoint Security Manager 1.6**, which lets network administrators centrally manage Comodo security products. The release offers a localized interface and an administration console that is now available in Chinese. Also, the release provides integration with Comodo Internet Security 4.1 and the ability to automatically install Comodo Internet Security in Chinese. Other additions to the release include compatibility with Windows Domain Controllers, and the ability to use custom ports for the antivirus component while safeguarding entry points to the user's environment. To learn more, visit www.comodo.com.

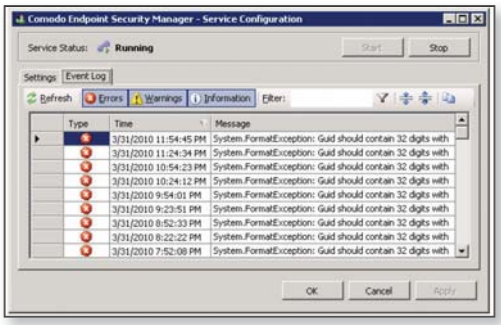

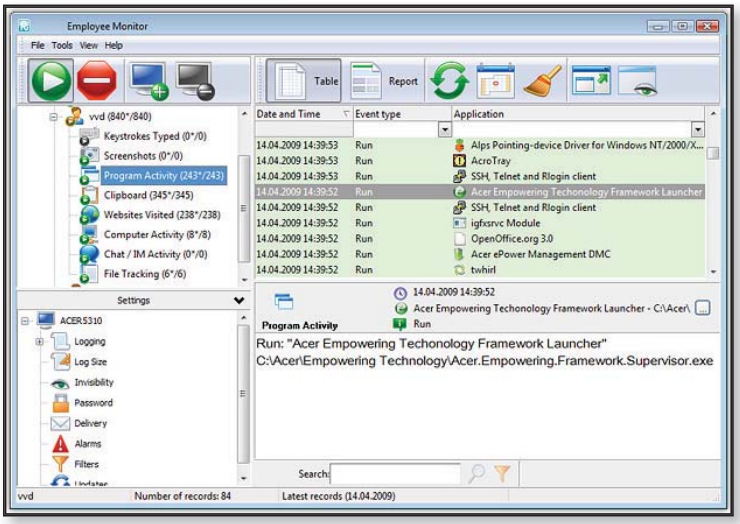

Paul's Picks

www.winsupersite.com

SUMMARIES of in-depth product reviews on Paul Thurrott's SuperSite for Windows

Small Business Server "Aurora" Release Candidate

PROS: Perfect starter server for small businesses; cross-premises foundation lets users pick the solutions that make sense for them

CONS: Use to start a new domain only—can't be added to existing environments

RATING: ◆◆◆◆◆

RECOMMENDATION: While Microsoft is hedging its bets with a more traditional, on-premises Windows Small Business Server 7 release this year, the company's "cross-premises" "Aurora" solution leaves a much more positive impression. Gone is the complexity of installing and maintaining multiple server solutions, replaced by a stripped-down but powerful Windows Server solution that offers best-of-breed storage and user management. Need email, calendaring, or other services? You're free to install on-premises servers as before, or utilize cloud services like Microsoft's Exchange Online (or even competitors like Google Apps), or you can mix and match. This is the right product at the right time for the small business market.

CONTACT: Microsoft • www.microsoft.com

DISCUSSION: www.winsupersite.com/server/aurora.asp

Xbox LIVE on Windows Phone

PROS: Gives Windows Phone a decisive and important advantage over all other smartphone platforms

CONS: Not all Xbox LIVE features are available on the phone; not all Windows Phone games will utilize Xbox LIVE

RATING: ◆◆◆◆◇

RECOMMENDATION: For months, the naysayers have been picking apart Windows Phone like carrion descending on road kill. "There's no cut and paste," they moan. "Where's the third-party multitasking?" they ask. Who cares? Windows Phone has Xbox LIVE, and this is the real deal, with gamercards, achievements, gamerscores, avatars, leader boards, multiplayer matchmaking, friends lists, and virtually everything else that makes Xbox LIVE special. It even has unique features like support for turn-by-turn games and a new "try before you buy" Trial Mode that game makers can easily add to their Windows Phone titles. On the phone, unlike the console, there's no need to pay for an Xbox LIVE Gold subscription. The inclusion of Xbox LIVE on Windows Phone, and an amazing selection of over 60 games at launch is going to put Microsoft's new smartphone system over the top. This really is a game changer.

CONTACT: Microsoft • www.microsoft.com

DISCUSSION: www.winsupersite.com/mobile/wp7_xbox_live.asp

InstantDoc ID 125859

REVIEW

SpamTitan

Messaging systems such as Microsoft Exchange Server 2010 are becoming more adept at dealing with spam and viruses, but generally they still need support from additional products. **SpamTitan** is a filtering product that runs on its own machine rather than on the messaging server; it sits in the message path to receive mail from the Internet and to forward mail from internal servers. It includes content filters based on MIME type and attachment type.

Installation and Setup

The installation process can be carried out either by downloading an ISO file and installing it on its own server or by downloading a pre-prepared virtual machine (VM) that runs on the VMware virtualization platform. For my testing, I downloaded the ISO file and installed it on a Microsoft Hyper-V VM. The hardware requirements for SpamTitan are very low. An old machine with a Pentium 4 CPU and 512MB of RAM should do just fine.

Installation is relatively straightforward, although dealing with the FreeBSD OS might be unfamiliar and feel clunky. However, if you read the documentation carefully, you shouldn't have a problem. The documentation is accurate, but it could be written or organized more clearly in places—for example, when explaining how to set up the disk partitions.

After installation, you must configure basic networking information and then define how the system should route mail. This process is easily done by following the system setup chapter of the administration guide, which should have you up and running in about half an hour. The documentation is adequate, although a little more detail in some areas—for example, the LDAP setup section—would have been nice, and on occasion it was necessary to jump to later chapters to fully understand the required settings.

SpamTitan uses two antivirus engines, ClamAV and Kaspersky, which throughout my testing caught all viruses they encountered. For spam filtering, SpamTitan uses many technologies, including Realtime Blackhole Lists (RBLs), whitelisting and blacklisting, a variety of email Request for Comments (RFC) compliance measures, and Bayesian analysis, which can all be enabled to support the built-in spam-filtering engine.

What's particularly useful about SpamTitan is the granular nature by which policies can be applied. You can configure different spam- and virus-filtering settings for different email domains and systems. You can also configure outbound disclaimers for each domain and customize the notifications received when an email message breaks a policy.

For high-availability setups, you can cluster two or more SpamTitan machines, with up to eight clustered nodes reported in production. However, clusters are supported only in the same site rather than in different geographic locations.

Performance and Reporting

SpamTitan works well for daily use. Administrators will find the web interface easy to use and well organized. You'll find that the online help, available through links on each dialog box, is more up-to-date than the PDF-based documentation.

Since I implemented SpamTitan, I've received very little spam—certainly less than 1 percent of mail received. Perhaps even more importantly, after I whitelisted a couple of mailing lists that occasionally had been blocked, I've had no other false positives. All viruses received have been correctly filtered.

SpamTitan has a range of options for notifying users and administrators of what's going on in the messaging system and with the mail that flows through it. For end users, notifications take the form of email messages detailing items that have been filtered and giving the option to delete, whitelist, or deliver those items.

The web interface lets users manage their personal whitelists and blacklists and search their personal quarantine area. When carrying out a quarantine search, users can narrow the results by sender, date range, and filter type. From the results, you can deliver, whitelist, or delete messages just as you can from a notification email message. One thing I found irritating was that although you can sort by spam score and date, you can't sort the results by sender address to find all instances of mail coming from a blocked mailing list.

Updates to virus and spam signatures are automatic; however, updates to the system software are not. In fact, you'll need to take care with this process because unless you have a clustered system or another way of maintaining mail flow, this procedure will require downtime.

Finally, the SpamTitan software includes a variety of reports with information such as top email recipients, top spam recipients, top viruses, and top spam relays. Reports are available manually but can also be scheduled for regular delivery to management. Even better, these reports can be output as a PDF or even in a spreadsheet to allow further data manipulation.

A Solid Choice

SpamTitan is quick and easy to set up and works well with limited configuration. It has good granularity of policies and a simple UI for both administrators and end users. I recommend SpamTitan as a cost-effective way to protect your organization from email-borne threats.

InstantDoc ID 125804

SpamTitan

PROS: Low system requirements; satisfactory spam filtering; granular policies

CONS: Documentation could be clearer; no Hyper-V VM to download

RATING: ◆◆◆◆◆

PRICE: $395 for 50 users for one year

RECOMMENDATION: SpamTitan is easy to set up, works well with limited configuration, and features granular policies. SpamTitan is a cost-effective way to protect your organization from email-borne threats.

CONTACT: SpamTitan • 201-984-3271 • www.spamtitan.com

Nathan Winters | nathan@clarinathan.co.uk

3 Disk-Imaging Solutions—Redux

A new look at three image-deployment products focuses on how their functionality differs

by Eric B. Rux

I first reviewed this comparative review's products—Acronis Snap Deploy, Paragon Deployment Manager, and Symantec Ghost—two and a half years ago, in "3 Disk-Imaging Solutions" (InstantDoc ID 98817). Although the market has changed, the basic concepts have remained the same. So, rather than rehash the same functionalities I discussed in that article, this time I want to take a different approach by investigating the features that set the products apart.

The goal of each product is to help you quickly deploy an OS to multiple computers. Instead of inserting a CD/DVD into each computer and running through the typical "Next, Next, Next" installation routine hundreds of times, you can use these imaging solutions to greatly streamline your deployment.

The Setup

Before I dive into how the individual products work, I'll summarize the task we want to accomplish: Essentially, we want to quickly and easily lay an OS down onto a fresh hard disk. All three products accomplish this task by taking a *disk image*, then giving you a way to distribute an exact copy of that image to a large number of computers. If you need to install Windows 7 or Windows XP onto just two or three computers, these solutions probably aren't for you; they'll take more time to set up and test than they're worth. But if you need to deploy ten, a hundred, or even a thousand computers, these products will help you immensely. These solutions are also useful for quick redeployment of machines in the wake of a virus infection, or after a computer has been reissued to a new user.

The first step in preparing to deploy an image is to create a *master image*. Start with a computer that best represents the kinds of systems that you have in your office. Modern versions of Windows do a good job of plug-and-play (PnP) driver installation for devices such as video, network, and sound cards, but they can struggle with major changes that affect the hardware abstraction layer (HAL), such as the type or number of CPUs. Keep this in mind when you're choosing the computer to represent your master image.

After installing the OS that you want to deploy and applying the latest service pack and patches, your next step really depends on your deployment philosophy. Some administrators deploy only the OS, whereas others deploy the OS as well as software such as Microsoft Office. Deploying only the OS allows for a very quick deployment, followed by a flexible installation of specific software via System Center Configuration Manager (SCCM), Group Policy, or another method. Deploying the image with both the OS and company software takes longer, but it's a simpler process because the software that users need is automatically installed. The choice is yours, but remember: The master image you create is immediately "outdated" the minute you create it—so, the less software you include, the better. I prefer to simply deploy the OS, then use Group Policy to deploy the software packages.

Finally, you need to run Sysprep. Doing so *generalizes* the computer and removes the computer name and the Security Identifier (SID). You're now ready to create the image, store it on a network share, and deploy it to computers. I'll tackle that discussion in my breakdowns of the three products. (Note that I installed all three products onto a standalone XP client that was not a member of a domain.)

Acronis Snap Deploy

Acronis offers a number of products for backup and disaster recovery of Exchange Server, SQL Server, and Small Business Server (SBS). Acronis Snap Deploy is targeted specifically to support personnel who need to mass-deploy client-based OSs. Figure 1 shows the Acronis Snap Deploy management console.

Getting the master image. Acronis Snap Deploy offers several ways to import an image of the master computer's hard disk. The Master Image Creation Wizard lets you create a bootable floppy, bootable CD/DVD, bootable ISO image, or bootable PXE configuration for the Acronis PXE Server. Regardless of the bootable media you choose, the result is the same: Take a cold snapshot of the master computer's hard disk and send it to a network share for later deployment. I tested the bootable ISO image (in a virtual machine—VM), as well as the bootable image via the included PXE server. Both worked flawlessly.

For some reason, the PXE server that comes with Acronis Snap Deploy requires a setup process. This configuration seemed unnecessary to me because the same settings are available when the bootable media is created (and then imported into the PXE server). This feature does provide for more granularity—but again,

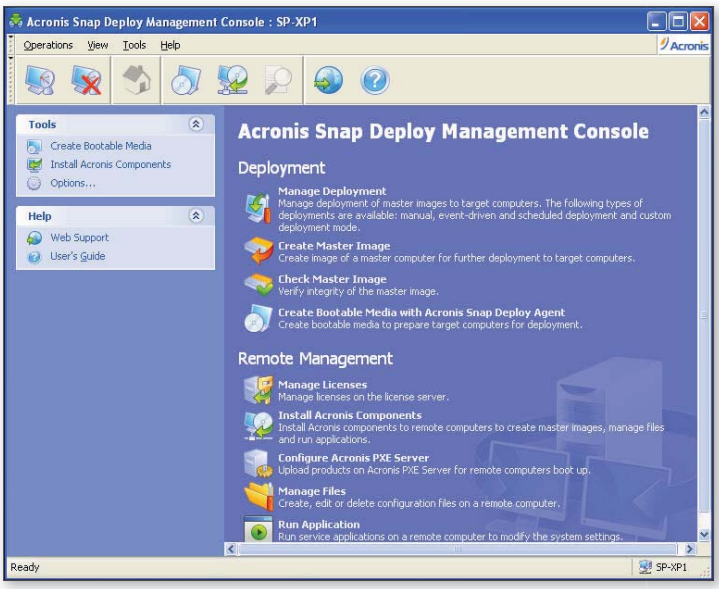

Figure 1: The Acronis Snap Deploy management console

it doesn't offer anything to the technician deploying the image; in other words, it should simply come preconfigured.

Deploying the master image. For simple one-computer deployments (e.g., following a virus attack), you can boot up the Acronis Snap Deploy Agent. You use the same boot media that you used to create the master image. A short wizard helps you find the deployment server and the master image you want to deploy.

However, if you have more of a project-type deployment (e.g., for a classroom full of 30 computers), you have to do a little more setup. First, you need to create the bootable media so that there's absolutely no user intervention required—again, this can be floppy, CD/DVD, PXE server, and so on. Click Create Bootable Media, and carefully choose the Acronis Snap Deploy Agent with an automatic start of 10 seconds. Doing so will cause the client to boot up the Acronis Snap Deploy Agent and immediately look for the deployment server for further instructions.

Those "further instructions" involve a configuration process via the Deployment and Templates tabs in the Manage Deployment section. First, select Event-Driven Deployment. This short wizard asks you two simple questions: *How many computers do you want to connect to the deployment server before it starts pushing out the image?* and *How many minutes/hours should the deployment server wait until it starts pushing out the image?* I appreciate that second question quite a bit. I've used other products that boast a "client count" feature only to have 1 computer out of 20 not cooperate—and I don't have a way to just push the image to the other 19. This feature lets you configure the product so that, after an hour, it will simply push the image out to the computers that have connected.

The second set of "further instructions" involves templates. Acronis templates let you specify machine-specific configuration. First, you select the master image you want to apply a template to. Then, you decide which physical disk to deploy to; determine whether to fit the image to the disk or create a partition; create an account on the target computer; determine a computer name; join a domain; set the IP address; determine whether you want to change the SID; choose files to copy to the target computer; and decide which—if any—applications to run after the deployment.

Extra features. In addition to its deployment features, Acronis Snap Deploy includes some basic remote management tools. You can create, edit, or delete files, and you can even start applications on remote computers. This functionality requires that you install an agent, so take that into consideration.

If you have a lot of computers with dissimilar hardware, you might want to consider an add-on called Acronis Universal Deploy. Whereas modern Windows versions do a good job of PnP identification for sound and video drivers, they don't handle major hardware differences such as the type or number of CPUs, motherboard brands, and so on. Acronis Universal Deploy lets you insert hardware-specific drivers into the image so that your master image is "universal" across all hardware types.

Acronis Snap Deploy

PROS: Simple to use; can define deployments by computer count or time elapsed

CONS: Some settings, such as the PXE server, seem unnecessary and confusing

RATING: ◆◆◆◆◇

PRICE: $25 per PC; see website for details

RECOMMENDATION: I recommend Acronis Snap Deploy for classroom or lab environments in which the maintenance of user data isn't a concern.

CONTACT: Acronis • www.acronis.com • 877-669-9749

Paragon Deployment Manager

Paragon Deployment Manager is similar in fit and function to the Acronis product, but it has some additional disk-utility features that set it apart from both of the other products. These features can be useful if you find yourself jumping back and forth between a disk-imaging solution and a disk-utility suite during your complex deployments. Figure 2 shows the Paragon Deployment Manager console.

Getting the master image. As soon as you finish installing Paragon Deployment Manager, you're ready to push an image of your master hard disk up to the network share. You have three boot options that will get you connected to Paragon Deployment Manager: Linux Boot CD, Windows PE Boot CD, and the included PXE server. The PXE server (a free, third-party application called Tftpd32) comes completely configured.

To capture the master image, you first need to map a drive to a network share (preferably the PC that's running Paragon Deployment Manager). To do so, select the Network Configurator and follow the wizard. Although this method worked in my tests, I found it a bit clunky when compared with Acronis's method. As soon as the client has a path to a network share, select Manual mode to start the copy of

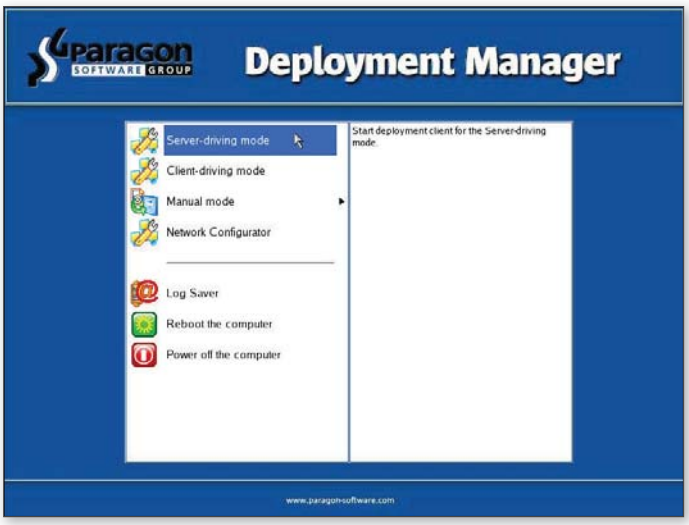

Figure 2: The Paragon Deployment Manager console

the hard drive to the network share. This process is very similar to that of Acronis Snap Deploy.

Deploying the master image. To deploy the master image to just one or two computers, you boot the target PC and manually point it to the server running Paragon Deployment Manager. For a more robust deployment strategy, you need to use Paragon Deployment Manager to create a new Session configuration. The Session can be automatic (in which the clients are set to deploy after they connect to the server), filtered by Session ID (the Session ID is set when the bootable media is created), or filtered by MAC address. One nice feature is the ability to add Post-Config options, such as reboot, power off, update files, and so on. Finally, you can schedule the deployment to recur on a specific schedule. This option could be useful for a weekly class that needs to have fresh computers every Monday morning, for example.

If you aren't necessarily running the reimage routines on a set schedule, but you need to repeat them from time to time, you'll want to create a Template. A Template defines a deployment the exact same way a Session does, but it lets you save the setting for later reuse.

Extra features. Paragon Deployment Manager also comes with Hard Drive Manager Professional. This handy utility lets you resize, merge, and undelete partitions; convert file systems; test the hard disk surface; and check the file-system integrity. There's even a tool to edit or view the individual hard disk sectors.

Paragon Deployment Manager

PROS: No PXE server setup required; can schedule deployments

CONS: Client mapping to image share requires clunky configuration

RATING: ◆◆◆◆◇

PRICE: $19 per PC; volume discounts available

RECOMMENDATION: I recommend Paragon Deployment Manager if you need to schedule image deployments so that the computers are fresh every Monday morning.

CONTACT: Paragon Software Group • www.deployment-manager.com • 888-347-5462

Symantec Ghost

Symantec Ghost differs from the other two products in this review. Whereas Acronis Snap Deploy and Paragon Deployment Manager are great at capturing an image of a hard disk and deploying it to multiple machines in cold environments (e.g., classrooms, labs), Ghost is designed to live with you and your users in production. The typical scenario in which Ghost shines involves a user who is running XP and needs to be upgraded to Windows 7. This functionality doesn't make Ghost better or worse than the other two products; it simply solves a different need. Figure 3 shows the Symantec Ghost console.

Getting the master image. Unlike the two previous products, Ghost doesn't use a PXE server. You *can* use a third-party PXE server if you want (Symantec even offers a detailed knowledge base article that can walk you through the setup process), but a PXE server isn't where Symantec has centered Ghost, so I'm not going to go there. Instead, all the tasks that you need to accomplish are managed from the Symantec Ghost Console.

To capture the master image, you first need to tell the Ghost server where to store the image (called the *image definition* in earlier versions). Next, you need to create a *Capture new image* task. This task tells Ghost which machine you want to capture and which drive or partition to include. You accomplish this entire process from the comfort of your desk; you never have to actually visit the computer you're capturing.

Deploying the master image. As in the previous step, you don't need to visit the computers that you're deploying an image to. Create a new task, and choose the operations that you want to perform. In this scenario, you would choose Clone, User Migration: Capture, and User Migration: Restore. Finally, choose the group of machines that you want to deploy to. Save the task, then click Execute. In a couple hours (depending on the amount of data

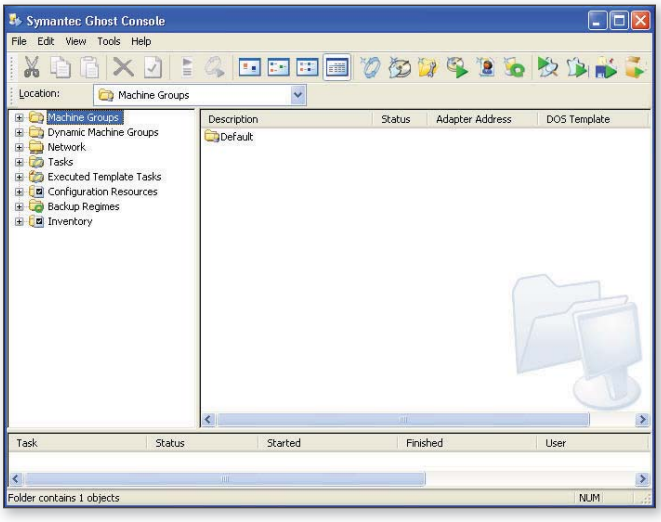

Figure 3: The Symantec Ghost console

on the user's computer), the user will be migrated and upgraded to the latest version of Windows.

Extra features. Ghost offers many more features than just the ability to deploy a new image of a hard disk. For starters, there's the user migration tool. If you've ever used XP's built-in Files and Settings Transfer Wizard, you understand the concept. Ghost lets you capture a user's files and settings, lay down a fresh image of a new OS, then restore those files and settings again.

Ghostcast Server lets you use Ghost more like Acronis Snap Deploy and Paragon Deployment Manager. It supports unicast, broadcast, and multicast deployments. In order for clients to connect to the Ghostcast Server, you'll need to get them booted up in a DOS-like environment (e.g., MS-DOS, PC-DOS) with network drivers for your particular network card. Ghost includes the Boot Wizard to help you through this process. Ghost supports PC-DOS, MS-DOS, Windows PE, and Linux.

Symantec Ghost

PROS: Can capture the user's data and settings, then put it all back after the new OS has been deployed

CONS: Cost

RATING: ◆◆◆◆◆

PRICE: $16.50–$66.26 per license; see website for details

RECOMMENDATION: If you need to do more than just deploy images, and you're tired of manual upgrades to the latest OS, you need to consider Symantec Ghost, which gets my highest recommendation.

CONTACT: Symantec • www.symantec.com • 800-745-6054

What Are Your Goals?

As always, before you make a purchase, be sure to review all the extra tools that come with these products. That being said, your choice of disk-imaging product will probably have more to do with your deployment goals than anything. For simple classroom or lab environments in which your only goal is to get a fresh image applied to a hard disk, Acronis Snap Deploy takes the cake for pure simplicity of deployment. Paragon Deployment Manager has some nice additional tools for manipulating hard disks that some users might find useful.

If you need to deploy OSs onto computers that already have existing users, Ghost is the clear choice. Users get to keep all the data that they forgot to store on the network, their desktop wallpaper stays intact, and you get to perform the entire migration from the comfort of your desk. Does it get much better than that?

InstantDoc ID 125797

Eric B. Rux
(ebrux@whshelp.com) is a contributing editor for *Windows IT Pro*, is cofounder of WHSHelp.com, and writes a monthly column at svconline.com/connectedhome. He is a senior Windows administrator and teaches the Microsoft Certified Systems Administrator (MCSA) program at a tech college.

Power Management Software for Windows Workstations

Save money by reducing workstations' energy consumption

by Karen Bemowski

GE is doing it, and Governor Arnold Schwarzenegger wants it done. "It" is reducing energy usage in IT operations. While Governor Schwarzenegger would like California's state-run IT operations to reduce their energy usage by 30 percent in the next two years (gov.ca.gov/press-release/14406), GE's IT managers are already doing so. By using Windows's power management features, they're saving GE more than $2.5 million annually (tinyurl.com/energystar-GE).

GE isn't the only one bringing good results to light. Companies such as FedEx, Dell, and Verizon have noted similar success stories. For example, FedEx is saving $1 million annually (tinyurl.com/energystar-FedEx).

Annual savings such as $2.5 million and $1 million go a long way to answer the question, "Why should IT implement power management policies on workstations?" There are other benefits as well, such as saving heat stress and wear on computers and monitors. Plus, there's the environmental benefit of reducing IT's carbon footprint.

There are many ways to reduce workstations' energy consumption using the power management settings built into Windows workstations. They include:

- Individually configuring each workstation's power management settings.
- Using Group Policy to centrally configure power management settings. This is possible in Windows 7 and Windows Vista workstations only.
- Using scripts with the Powercfg.exe command-line utility to centrally configure power management settings. Powercfg.exe is included with Windows XP SP2 and later.
- Using scripts with Windows Management Instrumentation's (WMI's) power policy classes to centrally configure power management settings. These WMI classes are available in Windows 7 only.

If you manage numerous workstations, have many older workstations, or don't want to write scripts, you might consider another option: using third-party power management software. This type of software typically lets you centrally configure and manage the power management settings for numerous workstations—including those running older client OSs—without writing any scripts. Plus, they often offer additional features such as reporting capabilities, enforcing compliance, and turning off workstations without using a utility such as Shutdown.exe.

This month's buyer's guide table gives you an overview of 12 power management programs for workstations. It includes both standalone power management software and power management software that's part of a larger suite.

Before you start looking for power management software, it helps to know some basics. Windows uses a combination of keyboard, mouse, and CPU activity to determine when a computer is idle. To ascertain the amount of time since the last activity, it uses a display idle timer and system idle timer. Windows compares the length of time noted by the idle timers to the settings specified through the Control Panel Power Options applet or a power management policy or plan. When the idle time is longer than the specified settings, the computer goes into the appropriate mode.

On desktop workstations, Windows OSs have four basic power saving modes: *System hibernate*, *System sleep* (*System standby* in Windows XP and earlier), *Turn off monitor*, and *Turn off hard disks*. According to ENERGY STAR, fully powered desktop PCs typically use about 65 watts and monitors use from 35 watts (LCDs) to 80 watts (CRTs). In hibernate mode and sleep mode, the PC and monitor use only 1 to 3 watts each. In *Turn off monitor* mode, the monitor uses 1 to 3 watts. Turning off the hard disks saves little energy.

You can customize the idle settings for these modes. The Environmental Protection Agency (EPA) recommends setting computers to enter sleep or hibernate mode after 30 to 60 minutes of inactivity. Lower settings can lead to more energy savings. The amount of money you'll actually save depends on many other factors as well, such as the workstations' current power management settings, the amount of time the workstations are left running, the amount of time they're actively used, their rate of power consumption, and the cost of electricity. There are online energy savings calculators you can use to get a ballpark idea of the energy savings (e.g., ComputersOff.org's calculator at www.computersoff.org/display.asp?id=17). Power management software that offers power savings reports can also give you an estimate of how much you can save.

POWER MANAGEMENT SOFTWARE

Shutting down workstations when they're not being used for long periods of time (e.g., nonbusiness hours) can save you the most money. Most power management programs let you schedule shutdowns in addition to scheduling the sleep or hiberate mode. Various technologies are used to power up workstations from hibernation or shutdown, including Wake on LAN (WOL), Wake on WAN, and Wake on Web (powering up a workstation from a remote location using a web browser).

Power management software can offer many other features, such as:

- Using additional metrics to determine when a computer is idle. For example, the software might analyze disk or application activity in addition to keyboard, mouse, and CPU activity.
- Preventing PC insomnia. PC insomnia occurs when legacy applications, open file handles, faulty mice, etc., prevent the machine from going into an idle

Company	Product	Price	Supported Windows Client OSs	
Adaptiva 425-823-4500 www.adaptiva.com	Adaptiva Green Planet (standalone software)	$8 to $12 per seat, depending on volume	Windows 7 (64-bit and 32-bit), Windows Vista (64-bit and 32-bit), Windows XP (64-bit and 32-bit), and Windows 2000 Professional	
Data Synergy 44 (0) 8456-435-035 www.datasynergy.co.uk	PowerMAN PC Power Management (standalone software)	Varies with quantity; about £5 ($7.75) per PC for a 2,000-seat installation	Windows 7 (64-bit and 32-bit), Vista (64-bit and 32-bit), XP (64-bit and 32-bit), and Win2K Pro	
EMCO Software 646-233-1163 www.emco.is	EMCO Remote Shutdown (standalone software)	From $99 (up to 25 PCs) to $389 (unlimited PCs in one location)	Windows 7 (64-bit and 32-bit), Vista (64-bit and 32-bit), XP (64-bit and 32-bit), and Win2K Pro	
EnviProt 49 7032 944 506 www.enviprot.com/en.html	Auto Shutdown Manager (standalone software)	Depends on license volume and target group; around $3 per PC for education institutions; less than $10 for lifetime licenses for key accounts; free light edition for private consumer use	Windows 7 (64-bit and 32-bit), Vista (64-bit and 32-bit), XP (64-bit and 32-bit), and Win2K Pro	
Faronics 604-637-3333 800-943-6422 www.faronics.com	Faronics Power Save (standalone software)	$12 per corporate license; $6 per education license	Windows 7 (64-bit and 32-bit), Vista (64-bit and 32-bit), XP (64-bit and 32-bit), and Win2K Pro	
Kaseya 415-694-5700 www.kaseya.com	Desktop Policy Management module (add-on for the Kaseya Enterprise Edition suite)	Suite: $17,000 (100-seat minimum); Desktop Policy Management module: $8 per seat	Windows 7 (64-bit and 32-bit), Vista (64-bit and 32-bit), XP (64-bit and 32-bit), Win2K Pro, and earlier	
Lakeside Software 248-686-1700 800-969-7717 www.lakesidesoftware.com	SysTrack Power Manager (part of the SysTrack suite)	Contact vendor	Windows 7 (64-bit and 32-bit), Vista (64-bit and 32-bit), XP (64-bit and 32-bit), Win2K Pro, and earlier	
NetWrix 201-490-8840 888-638-9749 www.netwrix.com	NetWrix Workstation Power Saving Manager (standalone software)	Starts at $5 per computer for 150 computers (as low as $0.40 per computer for more computers); freeware edition available for up to 50 workstations	Windows 7 (64-bit and 32-bit), Vista (64-bit and 32-bit), XP (64-bit and 32-bit), and Win2K Pro	
New Boundary Technologies 612-379-3805 800-747-4487 www.newboundary.com	PwrSmart (standalone software that can be on-premise or delivered as a service)	On-premise: $15 per workstation; monthly service fee based on number of subscribed PCs during the month	Windows 7 (64-bit and 32-bit), Vista (64-bit and 32-bit), XP (64-bit and 32-bit), and Win2K Pro	
ScriptLogic 800-813-6415 www.scriptlogic.com	Desktop Authority (suite)	$39 per seat with quantity discounts	Windows 7 (64-bit and 32-bit), Vista (64-bit and 32-bit), XP (64-bit and 32-bit), Win2K Pro, and earlier	
Sparxent's Verismic Software 44 (0) 1256-338-470 www.verismic.com	Verismic Power Manager (standalone software that can be on-premise, hosted, or delivered via Management Service Provider—MSP)	On-premise starts at £11 ($15) per node; hosted starts at £0.65 ($1.02) per node per month; contact vendor for MSP pricing	Windows 7 (64-bit and 32-bit), Vista (64-bit and 32-bit), and XP (32-bit)	
1E 866-592-4214 www.1e.com	NightWatchman (standalone software or part of the Power and Patch Management suite)	Contact vendor	Windows 7 (64-bit and 32-bit), Vista (64-bit and 32-bit), and XP (64-bit and 32-bit)	

state. Thus, the power management settings never kick in.
- Being aware of Windows Update settings.
- Auditing to determine compliance with power management policies.
- Blocking power management actions when certain applications are running.
- Automatically generating standard reports, such as power savings reports.

Before purchasing power management software, you should check with your power utility to see whether it offers any rebates. For example, Pacific Gas and Electric (PG&E) offers several different efficiency programs, including a $15 rebate for every networked PC that's licensed with qualified power management software (tinyurl.com/PGE-rebates).

InstantDoc ID 125800

Karen Bemowski
(kbemowski@windowsitpro.com) is a senior editor for *Windows IT Pro* and *SQL Server Magazine*.

Supported Windows Server OSs	Other Supported OSs	Manages the Power Settings on	Determines Inactivity Using	Customizable Idle Setting for			
				Sleep Mode	Hibernate Mode	Turn Off Monitor Mode	Turn Off Hard Disks Mode
Windows Server 2008 R2 (64-bit and 32-bit), Windows Server 2008 (64-bit and 32-bit), Windows Server 2003 R2 (64-bit and 32-bit), Windows Server 2003 (64-bit and 32-bit), and Windows 2000 Server		Desktop workstations, monitors, and laptops	Keyboard, mouse, CPU, and disk activity; memory utilization of running processes and applications	Yes	Yes	Yes	Yes
Server 2008 R2 (64-bit), Server 2008 (64-bit and 32-bit), Windows 2003 R2 (64-bit and 32-bit), and Windows 2003 (64-bit and 32-bit)		Desktop workstations, monitors, and laptops	Keyboard, mouse, CPU, and disk activity; specific programs running and specific files present	Yes	Yes	Yes	Yes
Server 2008 R2 (64-bit and 32-bit), Server 2008 (64-bit and 32-bit), Windows 2003 R2 (64-bit and 32-bit), Windows 2003 (64-bit and 32-bit), and Win2K		Desktop workstations					
Server 2008 R2 (64-bit and 32-bit), Server 2008 (64-bit and 32-bit), Windows 2003 R2 (64-bit and 32-bit), Windows 2003 (64-bit and 32-bit), and Win2K	For the idle analysis system, any TCP/IP-enabled device running MAC or Linux	Desktop workstations, monitors, laptops, servers, and virtual machines (VMs)	Keyboard, mouse, CPU, and disk activity; activity analysis of networks, processes, applications, printer queues, terminal sessions, busy/active phone lines, ongoing backups, updates, virus scanners, and more	Yes	Yes	Yes	
	Mac	Desktop workstations, monitors, and laptops	Keyboard, mouse, CPU, disk, and network activity	Yes	Yes	Yes	Yes
Server 2008 R2 (64-bit and 32-bit), Server 2008 (64-bit and 32-bit), Windows 2003 R2 (64-bit and 32-bit), Windows 2003 (64-bit and 32-bit), Win2K Server, and earlier	Mac and Linux (in beta)	Desktop workstations, monitors, laptops, and servers	Keyboard, mouse, CPU, and disk activity	Yes	Yes	Yes	Yes
Server 2008 R2 (64-bit and 32-bit), Server 2008 (64-bit and 32-bit), Windows 2003 R2 (64-bit), Windows 2003 (64-bit and 32-bit), Win2K, and earlier		Desktop workstations, monitors, laptops, printers, and VMs	Keyboard, mouse, CPU, disk, and other activity	Yes	Yes	Yes	Yes
Server 2008 R2 (64-bit and 32-bit), Server 2008 (64-bit and 32-bit), Windows 2003 R2 (64-bit and 32-bit), and Windows 2003 (64-bit and 32-bit)		Desktop workstations and laptops	CPU activity				
Server 2008 R2 (64-bit and 32-bit), Server 2008 (64-bit and 32-bit), Windows 2003 R2 (64-bit and 32-bit), and Windows 2003 (64-bit and 32-bit)		Desktop workstations, monitors, and laptops	Keyboard, mouse, CPU, and disk activity	Yes	Yes	Yes	Yes
Server 2008 R2 (64-bit and 32-bit), Server 2008 (64-bit and 32-bit), Windows 2003 R2 (64-bit and 32-bit), Windows 2003 (64-bit and 32-bit), Win2K Server, and earlier			Keyboard and mouse activity	Yes	Yes	Yes	Yes
Server 2008 R2 (64-bit and 32-bit), Server 2008 (64-bit and 32-bit), Windows 2003 R2 (32-bit), and Windows 2003 (32-bit)		Desktop workstations, monitors, and laptops	Keyboard, mouse, CPU, and running process activity	Yes	Yes	Yes	Yes
	Mac OS X	Desktop workstations, monitors, and laptops	Keyboard and mouse activity	Yes	Yes	Yes	Yes

POWER MANAGEMENT SOFTWARE

Company	Lets You Schedule					Powers Up Workstations Using	Centralized Control of Power Management (PM) Settings	PM Settings Applied Using	
	Sleep Mode	Hibernate Mode	Wake-Ups from Sleep or Hibernation	Shutdowns	Reboots				
Adaptiva 425-823-4500 www.adaptiva.com	Yes	Yes	Yes	Yes—Notifies users and saves open files before shutting down; users can temporarily exempt their computers from a shutdown	Yes	Wake on LAN (WOL), Wake on WAN, and Wake on Web	Yes	GPOs, registry settings, built-in technology (integrates with Microsoft System Center Configuration Manager—SCCM)	
Data Synergy 44 (0) 8456-435-035 www.datasynergy.co.uk	Yes	Yes	Yes	Yes—Notifies users before shutting down; users can temporarily exempt their computers from a shutdown	Yes	WOL, Wake on WAN, and Wake on Web	Yes	GPOs or registry settings	
EMCO Software 646-233-1163 www.emco.is	Yes	Yes	Yes	Yes—Notifies users before shutting down; users can temporarily exempt their computers from a shutdown	Yes	WOL		Built-in technology	
EnviProt 49 7032 944 506 www.enviprot.com/en.html	Yes	Yes	Yes	Yes—Notifies users and saves open files before shutting down; users can temporarily exempt their computers from a shutdown	Yes	WOL, Wake on WAN, Wake on Web, and Wake on WAN Client (users can wake up their office PCs remotely)	Yes	Built-in technology	
Faronics 604-637-3333 800-943-6422 www.faronics.com	Yes	Yes	Yes	Yes—Notifies users and saves open files before shutting down; users can temporarily exempt their computers from a shutdown	Yes	WOL	Yes	Built-in technology	
Kaseya 415-694-5700 www.kaseya.com	Yes	Yes	Yes	Yes	Yes	WOL, Wake on WAN, and Wake on Web	Yes	GPOs or built-in technology	
Lakeside Software 248-686-1700 800-969-7717 www.lakesidesoftware.com	Yes	Yes	Yes	Yes—Notifies users and saves open files before shutting down; users can temporarily exempt their computers from a shutdown	Yes	WOL, Wake on WAN, Wake on Web, and other mechanisms	Yes	Registry settings or built-in technology	
NetWrix 201-490-8840 888-638-9749 www.netwrix.com	Yes	Yes	Yes	Yes—Notifies users before shutting down; users can temporarily exempt their computers from a shutdown			Yes	GPOs, built-in technology, or configuration files	
New Boundary Technologies 612-379-3805 800-747-4487 www.newboundary.com	Yes	Yes	Yes	Yes—Notifies users before shutting down; users can temporarily exempt their computers from a shutdown		WOL and Wake on WAN	Yes	Built-in technology	
ScriptLogic 800-813-6415 www.scriptlogic.com	Yes	Yes	Yes	Yes—Notifies users before shutting down; users can temporarily exempt their computers from a shutdown	Yes	WOL and Wake on WAN	Yes	Registry settings and built-in technology	
Sparxent's Verismic Software 44 (0) 1256-338-470 www.verismic.com	Yes	Yes	Yes	Yes—Notifies users and saves open files before shutting down; users can temporarily exempt their computers from a shutdown	Yes	WOL, Wake on WAN, and Wake on Web	Yes	Built-in technology	
1E 866-592-4214 www.1e.com	Yes	Yes	Yes	Yes—Saves open files before shutting down; users can temporarily exempt their computers from a shutdown	Yes	WOL and timer-based wake-ups from sleep	Yes	GPOs, built-in technology, or command line	

Lets You				Has Separate PM Policy for When Nobody Is Logged On	Administrative Interface	Protects PM Policies from Unauthorized Changes	Is Aware of Windows Update Settings
Apply a PM Policy to an Existing AD Group	Define Groups of Workstations and Apply a PM Policy to Each Group	Define Groups of Users and Apply a PM Policy to Each Group	Apply More than One PM Policy to a Group				
Yes—Can exclude specific computers or users from AD groups	Yes	Yes	Yes		GUI	Yes	
Yes—Can exclude specific computers or users from AD groups	Yes	Yes	Yes	Yes	GUI	Yes	
Yes—Can exclude specific computers or users from AD groups	Yes		Yes		GUI		
	Yes	Yes	Yes	Yes	GUI and command line	Yes	
	Yes	Yes	Yes		GUI and command line	Yes	Yes
	Yes	Yes		Yes	GUI	Yes	Yes
Yes—Can exclude specific computers or users from AD groups	Yes	Yes	Yes	Yes	GUI and command line	Yes	Yes
Yes—Can exclude specific computers or users from AD groups		Yes	Yes		GUI	Yes	Yes
Yes—Can exclude specific computers or users from AD groups	Yes	Yes	Yes		GUI	Yes	
Yes—Can exclude specific computers or users from AD groups	Yes	Yes	Yes		GUI	Yes	
Yes—Can exclude specific computers or users from AD groups	Yes		Yes	Yes	GUI	Yes	
Yes	Yes				GUI and command line	Yes	Yes

POWER MANAGEMENT SOFTWARE

Company	Blocks PM Actions When Certain Applications Are Running	Prevents PC Insomnia	Auditing Capabilities	Automatically Generates Standard Reports	Customized or Interactive Reporting	Historical Reporting	Hosted Reporting	PM-Related Reporting Notes
Adaptiva 425-823-4500 www.adaptiva.com	Yes	Yes	Compliance auditing	Yes	Yes	Yes		
Data Synergy 44 (0) 8456-435-035 www.datasynergy.co.uk	Yes	Yes	Baseline and compliance auditing	Yes	Yes	Yes	Yes	One year of hosted reporting is included free, with a small fee thereafter; reporting is optional (client software works without it)
EMCO Software 646-233-1163 www.emco.is								
EnviProt 49 7032 944 506 www.enviprot.com/en.html	Yes	Yes		Yes	Yes	Yes	Yes	Can store all client events (e.g., shutdowns) in a central database for further analysis and reporting
Faronics 604-637-3333 800-943-6422 www.faronics.com	Yes	Yes	Baseline auditing					
Kaseya 415-694-5700 www.kaseya.com	Yes		Baseline and compliance auditing	Yes	Yes	Yes		
Lakeside Software 248-686-1700 800-969-7717 www.lakesidesoftware.com	Yes	Yes	Baseline and compliance auditing	Yes	Yes	Yes	Yes	Offers various reporting facilities
NetWrix 201-490-8840 888-638-9749 www.netwrix.com					Yes	Yes		Offers web-based reports
New Boundary Technologies 612-379-3805 800-747-4487 www.newboundary.com	Yes		Baseline auditing		Yes	Yes		
ScriptLogic 800-813-6415 www.scriptlogic.com								Provides a report of managed PCs and an estimation of energy savings for compliance purposes with power companies
Sparxent's Verismic Software 44 (0) 1256-338-470 www.verismic.com	Yes	Yes	Baseline and compliance auditing	Yes	Yes	Yes	Yes	
1E 866-592-4214 www.1e.com	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Management dashboard, baseline analysis reporting, and utility rebate reporting

Become an Exchange 2010 Maestro

A 3-Day Intensive Workshop

DATES & LOCATIONS

October 13–15, 2010 **Boston**

October 18–20, 2010 **Anaheim**

Windows IT Pro

Master this new technology with one-on-one help from the experts

Tony Redmond and Paul Robichaux

windowsitpro.com/go/PracticalExchange

You'll learn the key "gotchas" and hurdles that others have faced so that you don't have to!

- Best Exchange migration approaches
- Managing permissions
- What to know about Windows PowerShell scripts
- New compliance features
- Deploying Database Availability Groups

Worth your time and budget

Become an Exchange 2010 Maestro

windowsitpro.com/go/PracticalExchange

INSIGHTS FROM THE INDUSTRY

Multithreading, Multitasking PowerShell

A key characteristic of Windows PowerShell 1.0 was that it could basically only do one thing at a time. With 2.0, however, we get a brand new Background Jobs feature. There are three ways to start a job:

1. Use the -AsJob parameter of Get-WmiObject. This creates a new job with a default name, like "Job1" or "Job2."
2. Use the Start-Job cmdlet. This lets you specify a job name, and lets you create a job that runs one or more local commands or scripts.
3. Use the -AdJob parameter of Invoke-Command, causing your remote commands to run on a background thread. A -JobName parameter lets you specify a custom job name, if desired.

For example, let's say you run

```
Get-WmiObject -class Win32_BIOS
  -computer server,1server2,server3,
  server4,server5 -AsJob
```

When you create a new job, you're always creating a single, top-level master job and at least one child job. For jobs that connect to multiple computers, you'll have a child job for each computer. Start-Job only starts jobs that run on the local computer. The commands within that job can still connect to other computers, such as

```
Start-Job -command { Get-Service
  -computername server1,server2 }
```

In that case, it might be more efficient to use remoting to distribute the workload and communicate over WinRM. Compare the above command with the command

```
Invoke-Command -command { Get-Service
  } -computername server1,server2
  -asjob
```

Here, the remote computers are actually running Get-Service and returning their results to you, all over the single port required by WinRM.

You can run Get-Job to see what jobs are currently defined, running, completed, and so on; Stop-Job will kill a job that seems stuck, and Remove-Job will delete a job from memory. You can also use Wait-Job to cause the shell to pause until a specified job has completed; this can be useful inside of a script if you need to start a job and then have the script wait for it to finish. All of these job cmdlets accept parameters like -id or -name that let you specify which job or jobs you want to manage.

Here's a trick: Say you have a job, Job1, which is connecting to four computers. It will contain child jobs, one for each computer. To see them, run

```
Get-Job -name Job1 | Format-List *
```

You'll see the property that lists the child jobs' names, which might be Job2, Job3, Job4, and Job5. You can use those job names to manage just one of the child jobs. Note that the master job will only show "Completed" as its status if *every* child job also completed; the failure of a single child will result in a "Failed" status for the master job. Therefore, it's pretty useful to be able to check on the individual child jobs.

When there are results for a job, you get them by using the Receive-Job cmdlet. Receiving results for a master job grabs all of the child jobs' results; you can also specify a -name or -id to get the results for a child. Results are buffered in memory until they are received by you. Keep in mind that jobs are one of PowerShell's extensibility points. Other developers can write cmdlets that create different kinds of jobs.

—Don Jones

vSphere 4.1: More Virtualization Bang for your SMB Buck

Wow. I leave for a week's worth of unplugged vacation, and the VMware world changes (at time of this writing back in mid-July). We've known that vSphere 4.1 has been coming for some time, but its publicly-available details have been slim at best. Now that has changed.

I've long been a fan of VMware's products, yet until recently not the biggest fan of their pricing. Particularly at its lower-end price points, VMware's flagship product prior to this release offered too little for too much. This update swings the capabilities we've come to think of as "must have" well into the court of affordability for small and medium businesses.

I can think of four new virtualization bangs that ease the strain on the SMB buck:

VMware vMotion with the VMware vSphere 4.1 Essentials Plus and Standard editions. If you're a small or medium business, trying to get by on as small a budget as possible, this inclusion elevates your virtual environment far closer to those of the big boys. It adds the layer of high availability that you need to protect yourself and your virtual workloads.

VMware's public announcement that ESXi is the wave of the future. We've heard whispers of this narrative for years now, but without the absolute decision that we get with 4.1. In comparison to ESX, "i" is for SMBs a smaller package that can be less complex and more embraceable, while still being the fully-featured powerhouse that enterprises require.

Improvements to VMware Data Recovery (VDR). SMBs and SMB budgets need simple and cost-effective solutions that don't require them to look elsewhere for additional software. If this version's VDR becomes the backup solution we all wish it to be, SMBs will absolutely benefit from its inclusion in Essentials.

The simple fact that Essentials is less expensive than before. In vSphere 4.1, Essentials can go as low as $495, a great price point.

Huzzah to VMware for throwing a bone to SMB pocketbooks, while at the same time offering the new technologies that enterprises demand.

Note: In case VMware's product names get you confused: vMotion is now available in Essentials Plus and vSphere Standard. If you purchase the least-expensive Essentials kit, you won't get vMotion.

To read more virtualization tips and tricks, check out my blog at www.windowsitpro.com/blogs/Virtualization-ProTips.aspx.

—Greg Shields

Windows IT Pro Twitter Feeds to Follow

@windowsitpro - Get the latest article updates across the Windows IT Pro website.

@sqlservermag - Updates and giveaways from our sister publication, SQL Server Magazine.

@savvyasst - Related events, resources, and giveaways in the Windows IT space.

@michelecrockett - Updates from Michele Crockett, Editorial and Custom Strategy Director

@witproAmy - Updates from Amy Eisenberg, Executive Editor

@wincaroline - Updates from Caroline Marwitz, editor specializing in Active Directory and SharePoint

@breinholz - Updates from Brian Reinholz, editor specializing in training, certification, and mobility

@zacwiggy - Updates from Zac Wiggy, editor specializing in systems management, Windows OS, and virtualization

@bkwins - Updates from B. K. Winstead, editor specializing in Exchange, Outlook, and mobility

@thurrott - News Editor Paul Thurrott's Twitter feed

JOIN THE MOBILE REVOLUTION WITH THE WOR

WEBSITES

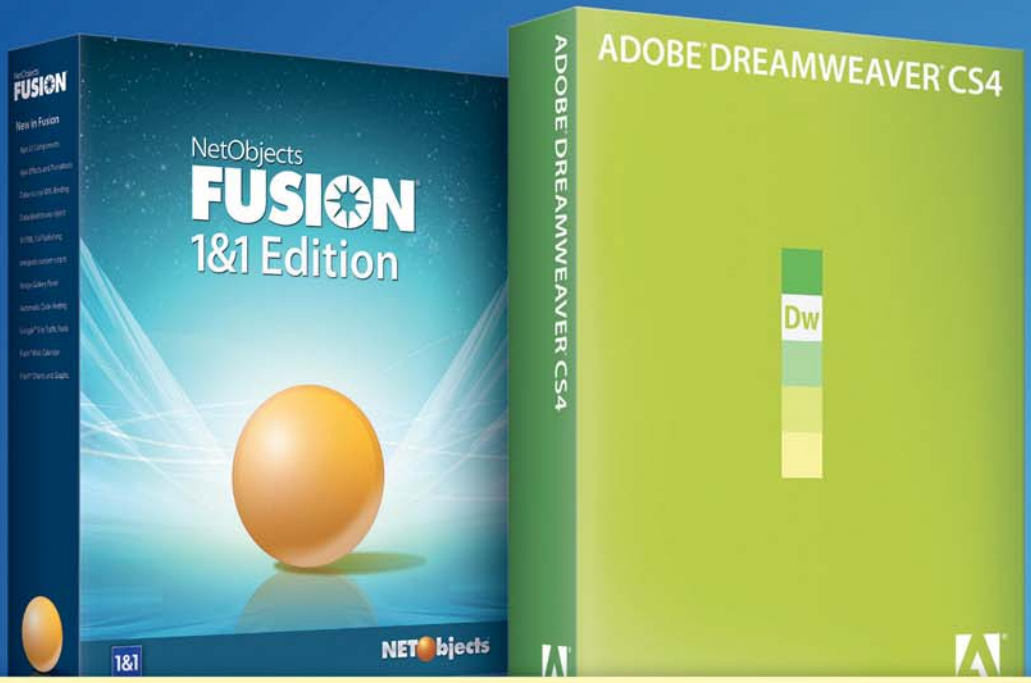

FREE SOFTWARE*

INCLUDED WITH 1&1 HOSTING PLANS. CHOOSE FROM:

NetObjects Fusion® 1&1 Edition is a website design application which creates sites that are optimized for mobile viewing. The 1&1 Edition was designed specifically for 1&1 web hosting packages and includes additional mobile templates as an extra bonus.

Adobe® Dreamweaver® CS4 is a sophisticated website design application for creating professional websites. Dreamweaver® includes the Adobe® Device Central module, enabling web designers to test their websites on mobile devices by emulating the latest smartphones.

TURN PAGE FOR DETAILS OR VISIT WWW.1AND1.COM.

EPA GREEN POWER PARTNER

LD'S LARGEST WEB HOST. AT 1&1 INTERNET:

GO MOBILE

More and more people are browsing the web via their iPhone, BlackBerry®, or smartphone. Don't miss out on these customers! At 1&1, you get the software you need to create additional websites that are optimized for mobile viewing.

- ✓ **Layouts, designs and wizards enhanced for the latest smartphones, like iPhone and BlackBerry®**
- ✓ **Compatible across multiple platforms**
- ✓ **Valued at up to $479!**

1&1®

MEMBER OF united internet

Get started today, call 1-877-GO-1AND1 **www.1and1.com**

*Software offer valid with select 1&1 web hosting plans, and is available for download in the 1&1 Control Panel only. 12 month minimum contract term, setup fee, and other terms and conditions may apply. Visit www.1and1.com for full promotional offer details. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 Internet AG, all other trademarks are the property of their respective owners.

Derelict Administrator Accounts: A Millennium Falcon Problem

Many system administrators have the same attitude toward the networks they manage that Han Solo has towards the Millennium Falcon. The cardinal rule is, "If it is currently working, don't mess with it."

That's why Han Solo got angry with Chewbacca for performing preventative maintenance in the Rebel Hangar on Hoth. The ship was working and then Chewie started messing with it. Han knew that pulling on any one thread could unravel the whole kit and caboodle.

There are a whole lot of loose threads that hang out about a network that it is tempting to tug on. One such thread that many administrators are reluctant to pull on is removing the user accounts of systems administrators who no longer work at the organization.

The main reason that people are reluctant to do anything about these accounts is a fear that if they disable the account, something—a script, a service, or something else in the entrails of the network infrastructure—will break. Better to let sleeping dogs lie, to not pull on a thread that may unravel more trouble than it is worth.

It is the Millennium Falcon problem. Start working on the landing gear and suddenly the hyper drive doesn't work. We've all had a bad experience when maintaining a network where we have started doing some routine maintenance on one thing, only to have something else that seems unrelated fail spectacularly. And let's face it: Most system administrators have enough fires to put out without worrying about pulling on threads that might start more.

So what can you do about the derelict accounts of former system administrators? Audit them. If a domain administrator account is being used to support a script or service, it has to be logging on. You can run a query from Active Directory Users and Computers to figure out which accounts haven't logged on recently. If you have someone who left more than a year ago but their account isn't on the list of accounts that haven't logged on for more than 30 days, you've certainly got an issue that you should investigate. If the account is on the list of accounts that haven't logged on for more than 30 days, then you can be a little more confident that disabling the account (with a view to eventual deletion) is unlikely to break the hyper drive. Read more on my blog at www.windowsitpro.com/blogs/hyperboleembellishmentsysadmin.aspx.

—Orin Thomas

Unified Communications in 2010: Where Are We with UC?

When it comes to unified communications (UC), I'm not entirely sure everyone is exactly on the bandwagon yet—at least not wholeheartedly. It seems that most organizations that claim to have implemented UC have done so only in a very limited way. For example, in my company, I get my voicemail messages and missed call notices in my Outlook Inbox, which is made possible by the unified messaging (UM) capabilities of Microsoft Exchange Server 2007.

But the potential for UC goes well beyond that. If you've seen any of the info on Office 2010 and its integration with SharePoint 2010, you might begin to get an idea. For instance, multiple people can check out the same document from SharePoint and work on it at the same time on their local computers; you can see who's working on a particular section at the moment and start an IM conversation with them right from within Word. Meanwhile, SharePoint keeps you from overwriting changes someone else is making and combines all updates into a newly revised document.

This example is, of course, a Microsoft-centric application of UC, but the potential is there to use that same presence information in the development and implementation of third-party applications or custom, in-house applications as well. Up to this point, most UC has been focused on the communications channel itself—unifying email, phones, conferencing, and so forth—because that's the easiest to implement and it seems like a big, flashy win for end users. Going to that next level of integrating communications throughout the line-of-business applications stack is going to be a lot more difficult to implement, and therefore a lot more difficult for companies and their overworked IT departments to embrace—but it's also the area with the potential for the greatest productivity gains and ROI.

It seems that most organizations that have implemented UC have done so only in a very limited way.

Naturally, I'd like to hear other views on the state of UC out there in the wild. Visit the Exchange & Outlook blog at www.windowsitpro.com/blogs/exchange-andoutlook.aspx to read similar articles and get back to me. Or send an email to bwinstead@windowsitpro.com explaining what you're doing, what headaches you've had to overcome, what benefits UC has presented to you and your business, or what's keeping you from implementing UC in the first place.

—B. K. Winstead

BIG SAVINGS

ON PROFESSIONAL WEB SOLUTIONS

With 1&1, you get premium web design software, and 50% off the first 6 months on our most popular web hosting plans.

NOW GET 50% OFF PLUS FREE SOFTWARE*

1&1® HOME PACKAGE

- 2 Domain Names Included (.com, .net, .org, .info or .biz.)
- 150 GB Web Space
- UNLIMITED Traffic
- 10 FTP Accounts
- 25 MySQL Databases
- Extensive Programming Language Support: Perl, Python, PHP4, PHP5, PHP6 (beta) with Zend® Framework
- NetObjects Fusion® 1&1 Edition

~~$6.99 per month~~

$3.49* per month for the first 6 months

1&1® BUSINESS PACKAGE

- 3 Domain Names Included (.com, .net, .org, .info or .biz.)
- 250 GB Web Space
- UNLIMITED Traffic
- 25 FTP Accounts
- 50 MySQL Databases
- Extensive Programming Language Support: Perl, Python, PHP4, PHP5, PHP6 (beta) with Zend® Framework
- NetObjects Fusion® 1&1 Edition or Adobe® Dreamweaver CS4

~~$9.99 per month~~

$4.99* per month for the first 6 months

1&1® DEVELOPER PACKAGE

- 5 Domain Names Included (.com, .net, .org, .info or .biz.)
- 300 GB Web Space
- UNLIMITED Traffic
- 50 FTP Accounts
- 100 MySQL Databases
- Extensive Programming Language Support: Perl, Python, PHP4, PHP5, PHP6 (beta) with Zend® Framework
- NetObjects Fusion® 1&1 Edition or Adobe® Dreamweaver CS4
- NEW: 1&1 Power Plus Performance Guarantee

~~$19.99 per month~~

$9.99* per month for the first 6 months

ALSO ON SALE:

.us domains $0.99/first year*

.com domains $7.99/first year*

Visit our website for a full list of special offers.

1&1®

MEMBER OF united internet

Get started today, call 1-877-GO-1AND1 **www.1and1.com**

*12 month minimum contract term required for software offer. Setup fee and other terms and conditions may apply. Software available for download in the 1&1 Control Panel only. Domain offer valid first year only. After first year, standard pricing applies. Visit www.1and1.com for full promotional offer details. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 Internet AG, all other trademarks are the property of their respective owners.

AD INDEX

For detailed information about products in this issue of *Windows IT Pro*, visit the web sites listed below.

COMPANY/URL	PAGE	COMPANY/URL	PAGE	COMPANY/URL	PAGE
1&1 Internet www.1and1.com	74, 75, 77	**IBM Corporation** www.ibm.com/questions	25	**ScriptLogic Corporation** www.scriptlogic.com/didyouknow	Cover Tip
Diskeeper Corporation www.diskeeper.com/v2	2	**Jalasoft Inc.** www.jalasoft.com/wings2010	79	**Sunbelt Software Inc.** www.TestDriveVipre.com	Cover 3
HP www.hp.com/servers/unleash12	Cover 4	**NetWrix Corporation** www.netwrix.com	4	**Train Signal** www.trainsignal.com	8
HOB Inc. www.hobsoft.com	27	**PowerWF Studio.** www.powerwf.com/mg1	72	**Vision Solutions Inc** www.visionsolutions.com	13
IBM Corporation www.ibm.com/systems/x3690	Cover 2	**Privacyware** www.privacyware.com	73	**WinConnections Fall Event** www.WinConnections.com	52, 53
IBM Corporation www.lotusknows.com	11	**Quest Software.** www.quest.com/liberating	37	**Windows IT Pro** www.windowsitpro.com	7, 64, 71

VENDOR DIRECTORY

The following vendors or their products are mentioned in this issue of *Windows IT Pro* on the pages listed below.

1E....66
3PAR....58
Acronis....61
Adaptiva....66
Centrify Corporation....58
Comodo....59
Data Synergy UK Ltd....66
Dell KACE....58
Diskeeper Corporation....58
EMCO Software....66
EnviProt....66
Faronics....66
Kaseya....66
Lakeside Software....66
NetWrix....66
New Boundary Technologies....66
Paragon Software Group....62
Red Gate Software....58
REFOG....59
ScriptLogic....66
SpamTitan....60
Sparxent's Verismic Software....66
Symantec....63
Symplified....58
VMware....73

DIRECTORY OF SERVICES | WINDOWS IT PRO NETWORK

Search our network of sites dedicated to hands-on technical information for IT professionals.
www.windowsitpro.com

Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
www.windowsitpro.com/go/forums

News

Check out the current news and information about Microsoft Windows technologies.
www.windowsitpro.com/go/news

EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

- *asp.netNOW*
- *DevProConnections UPDATE*
- *Exchange & Outlook UPDATE*
- *Security UPDATE*
- *SharepointPro Connections UPDATE*
- *SQL Server Magazine UPDATE*
- *Windows IT Pro UPDATE*
- *Windows Tips & Tricks UPDATE*
- *WinInfo Daily UPDATE*

www.windowsitpro.com/email

RELATED PRODUCTS

Custom Reprint Services

Order reprints of *Windows IT Pro* articles. Diane Madzelonka at Diane.madzelonka@penton.com.

Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the Web. Includes FREE access to eBooks and archived eLearning events, plus a subscription to either Windows IT Pro or SQL Server Magazine.
www.windowsitpro.com/go/vipsub

SQL SERVER MAGAZINE

Explore the hottest new features of SQL Server, and discover practical tips and tools.
www.sqlmag.com

ASSOCIATED WEBSITES

DevProConnections

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
www.devproconnections.com

SharePointPro Connections

Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and Web seminars mentored by a community of peers and professionals.
www.sharepointproconnections.com

NEW WAYS TO REACH *WINDOWS IT PRO* EDITORS:

LinkedIn: To check out the *Windows IT Pro* group on LinkedIn, sign in on the LinkedIn homepage (www.linkedin.com), select the Search Groups option from the pull-down menu, and use "Windows IT Pro" as your search term.

Facebook: We've created a page on Facebook for *Windows IT Pro*, which you can access at: http://tinyurl.com/d5bquf. Visit our Facebook page to read the latest reader comments, see links to our latest web content, browse our classic cover gallery, and participate in our Facebook discussion board.

Twitter: Visit the *Windows IT Pro* Twitter page at www.twitter.com/windowsitpro.

Windows IT Pro

Xian
Wings 2010

For BlackBerry®, Windows Mobile® and soon for iPhone®

System Center Operations Manager. Anytime. Anywhere.

Going Mobile? Take Microsoft® System Center Operations Manager with you.

Just another day out of the office. You're trying to watch a movie but your servers have different plans. One machine is complaining about drive space and another is having problems talking to the database. Both need to be fixed right away. Wouldn't it be nice to keep control of your IT systems anywhere and anytime you want?

With Xian Wings 2010 on your smartphone you can:

- Receive event notifications
- Review system and application status
- Build graphs for key performance data
- Run custom tasks
- Manage alerts

Get a free trial now!
here: www.jalasoft.com/wings2010

1219 Durham | Houston, Texas 77007 | Toll Free: 888 402-6717 |

Jalasoft

by Jason Bovberg

Sedentary No More!

PRODUCT OF THE MONTH

Here at the offices of *Windows IT Pro*—and at businesses around the world—we could surely benefit from a little more time on our feet. We spend far too much time in a sedentary position, just typing away. According to TrekDesk, "Medical reports and studies focusing on rising obesity rates and ill health caused by 'Sitting Disease' (a term given to the multitude of syndromes and diseases caused by sedentary lives) have emerged from numerous medical authorities this past year." In that light, our healthy curiosity was piqued by the TrekDesk Treadmill Desk, which the company touts as "an invaluable method to restore health and lose weight" at work. *Forbes* recently praised the TrekDesk Treadmill Desk as "one of the best workplace luxuries anywhere." At $479—"approximately one third the cost of a monthly health insurance premium," says TrekDesk—the TrekDesk Treadmill Desk strikes us as something that could change America. And for your company's next cost-cutting measure: All workers use such treadmills to generate electricity for their computers. Now, that's what we call green computing! Find more information about the TrekDesk Treadmill Desk at www.trekdesk.com.

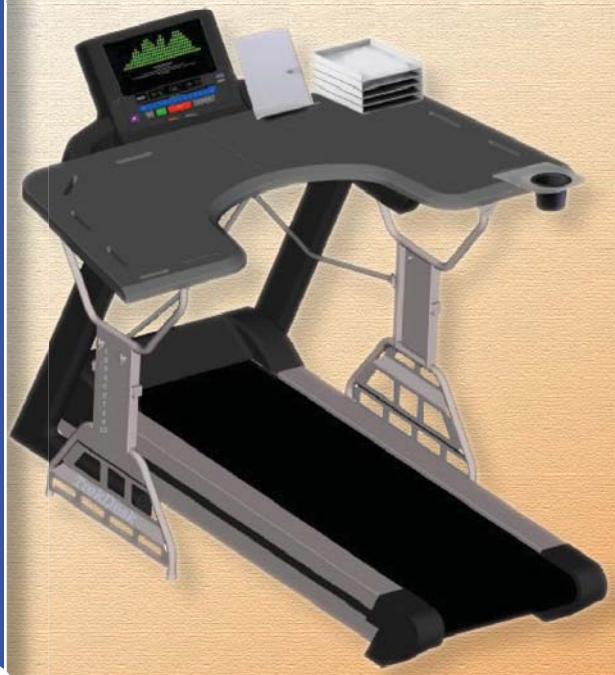

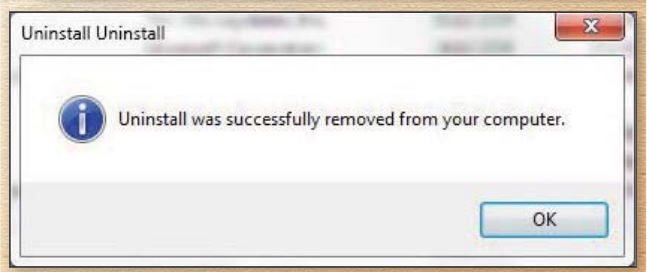

Figure 1: Will that create a black hole?

Figure 2: Nah, I don't need that

User Moment of the Month

In the early 1990s, my company purchased lunchbox portable computers for our sales staff to provide in-the-field demos to potential clients. These computers used the old floppy disks to load programs and upgrades. We sent out the floppies with detailed installation instructions. One particular salesperson called for assistance, saying, "I followed the instructions step by step and couldn't complete the upgrade." None of the other sales staff had experienced problems. I asked him, "At what point did you run into a problem?" He replied, "I loaded the first four disks with no problem, but the fifth disk won't fit in the slot." I was puzzled, but then asked, "You did remove each floppy after it ran, right?" He said, "No, the instructions said to load the five floppy disks. They didn't say to load them one at a time!"

—Cathie Crawford

October 2010 issue no. 194, *Windows IT Pro* (ISSN 1552-3136) is published monthly. *Windows IT Pro,* 221 E. 29th St., Loveland, CO 80538, (800) 793-5697 or (970) 663-4700. Sales and Marketing Offices: 221 E. 29th St., Loveland, CO 80538. Advertising rates furnished upon request. Periodicals Class postage paid at Loveland, Colorado, and additional mailing offices. **POSTMASTER: Send address changes to *Windows IT Pro,* 221 E. 29th St., Loveland, CO 80538. SUBSCRIBERS:** Send all inquiries, payments, and address changes to *Windows IT Pro,* Circulation Department, 221 E. 29th St., Loveland, CO 80538. Printed in the USA.

Kiss your antivirus bloatware goodbye

NEW Premium Version with Firewall

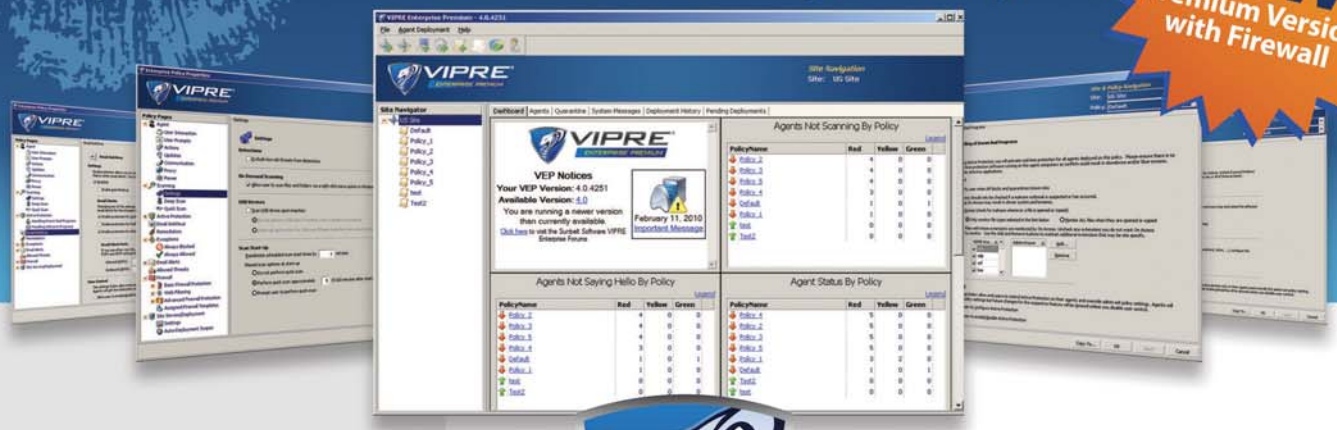

VIPRE®

TEST DRIVE ENTERPRISE PREMIUM

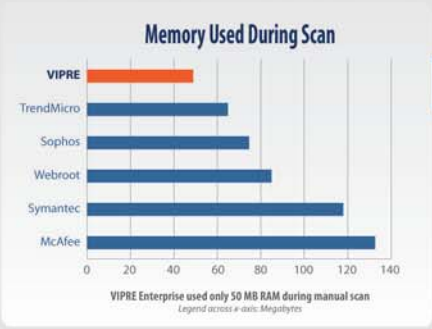

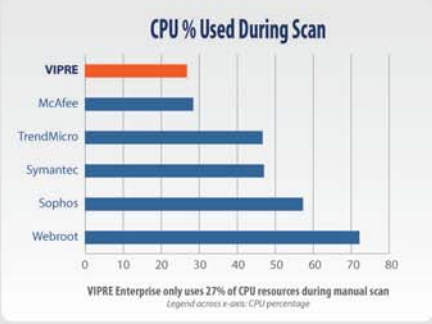

How does your current software compare? VIPRE Enterprise scans at a brisk 13.95 MB/sec and uses just 27% of CPU and 50 MB of RAM. In idle, it uses a mere 13.3 MB RAM with a disk footprint of just 113 MB. You'll hardly notice it's running!

Special Competitive Upgrade: *50% Discount!*

Until now, antivirus engines have been Frankensteins, bolted together from bits and pieces of different products. They're slow, full of bugs, and hard to manage.

VIPRE Enterprise Premium is a revolutionary new approach. It combines high-performance antivirus, antispyware, and desktop firewall into a single agent so you get comprehensive endpoint malware protection with low system resource usage. It's fast, powerful and easy.

Plus, advanced anti-malware technology protects your system against the new wave of malware threats. No more juggling multiple programs. No more dealing with user complaints about slow workstation performance.

- **COMPLETE!** All-in-one protection from today's malware.
- **FAST!** High-performance and low impact on system resources.
- **EASY!** Manage everything easily from one command screen.
- **RELIABLE!** Configurable, real-time monitoring technology.
- **AFFORDABLE!** Ask for a quote with our 50% competitive upgrade discount!

Why struggle with slow resource hogs when you can manage ALL your malware threats with one fast, easy application?

Curious? Download your FREE copy of VIPRE Enterprise Premium and give it a test drive.

When you compare VIPRE Enterprise Premium to Symantec, McAfee, Trend Micro or whatever antivirus program you're using, **you WILL want to switch!** Don't worry, though. You can get VIPRE Enterprise Premium with a **50% competitive upgrade discount!**

Compatible with Windows 7

BEST IN TECH 2009 Windows IT Administrator

Windows IT Pro 2009 Community Choice GOLD

vb 100 VIRUS virusbtn.com

Sunbelt Software
Part of the GFI Software Family

Plus we will buy out your existing maintenance contract for 1 year!

Download now: www.TestDriveVipre.com

Sunbelt Software Tel: 1-888-688-8457 or 1-727-562-0101 Fax: 1-727-562-5199 www.SunbeltSoftware.com sales@sunbeltsoftware.com

AMD Opteron™ 64

UNLEASH

faster server ROI.

Servers that pay for themselves in as little as 2 months: next generation HP ProLiant servers powered by AMD Opteron™ 6100 Series processors.[1]

Is an aging IT infrastructure costing you money? Now is the time to unleash the full potential of your business with next generation HP ProLiant servers powered by AMD Opteron™ 6100 Series processors:

- 23 to 1 server consolidation ratio[1]
- 96% or more energy and cooling savings[1]
- $48,380 saved per 100 users[2]

Accelerate your business and lay the groundwork for the HP Converged Infrastructure, the road map to greater IT efficiency. And unleash faster server ROI today.

Outcomes that matter.

HP ProLiant DL385 G7 Server

- AMD Opteron™ Processor Model 6134
- 4 GB memory, up to 256 GB Max
- Up to 8 small form factor high-performance SAS hard drives with standard cage. Or up to 16 SFF or 6 LFF hard drives with optional drive cages.
- Integrated Lights-Out 3 (iLO 3) providing industry-leading management and powerful administration

$2,599 (Save $498)

Lease for just $69/mo.*

Smart Buy (PN: 605869-005)

Calculate your ROI and register for *The Time is Right to Transform the Data Center* and *The Next Generation HP ProLiant Server Line: A Powerful Platform for Virtualization* white papers at **hp.com/servers/unleash12**

20 YEARS
OF x86 SERVER INNOVATION

*Prices shown are HP Direct prices; reseller and retail prices may vary. Prices shown are subject to change and do not include applicable state and local taxes or shipping to recipient's address. Offers cannot be combined with any other offer or discount and are good while supplies last. All featured offers available in U.S. only. Savings based on HP published list price of configure-to-order equivalent (DL Server: $3,097-$498 instant savings = Smart Buy price of $2,599.) Financing available through Hewlett-Packard Financial Services Company and its subsidiaries (HPFSC) to qualified commercial customers in the U.S. and is subject to credit approval and execution of standard HPFSC documentation. Prices shown are based on a lease 48 months in term with a fair market value purchase option at the end of the term and are valid through July 31, 2010. Other rates apply for other terms and transaction sizes. Financing is available on transactions greater than $349. Other charges and restrictions may apply. HPFSC reserves the right to change or cancel this program at any time without notice. This offer cannot be combined with any other rebate, discount or promotion without prior approval by HP and HPFSC. Rates are based on customer's credit rating, financing terms, offering types, equipment type and options. Not all customers may qualify for these rates. Other restrictions may apply. HPFSC reserves the right to change or cancel this program at any time without notice.

[1] Based on HP internal testing comparing hardware on HP ProLiant DL380 G4 to HP ProLiant DL385 G7.

[2] Source: IDC white paper sponsored by HP *Gaining Business Value and ROI with HP Insight Control*, May 2009.

hp